Particle Accelerators Find a New Role › Treating cancer
P.32

Keeping Power-Line Workers Safe › Better gear wards off deadly current
P.40

Chess by Telegraph › Soon after the first line opened, a telechess match commenced
P.42

FOR THE
TECHNOLOGY
INSIDER
MARCH 2026

IEEE Spectrum

AI Goes Hunting for Physics That Shouldn't Exist

IEEE

LIVE
LIFE ON
YOUR
TERMS.
INSURANCE
THAT FITS
YOU.
Group Term Life Insurance
engineered for you.

1-800-493-IEEE (4333)

To learn more*, visit IEEEinsurance.com/TLplan

IEEE
LIFE INSURANCE

*For information on features, costs, eligibility, renewability, limitations and exclusions.

Group Term Life Insurance is available only for residents of the U.S. (except territories), Puerto Rico and Canada (except Quebec). This is underwritten by New York Life Insurance Company, 51 Madison Ave., New York, NY 10010 on Policy Form GMR. AMBA does not act as a broker with respect to Canadian residents and acts solely as an Administrator on behalf of New York Life.

Association Member Benefits Advisors, LLC (AMBA).
In CA d/b/a Association Member Benefits & Insurance Agency
CA License #0I96562 • AR License #100114462
Program Administered by AMBA Administrators, Inc.

107580 (1/26) Copyright 2026 AMBA. All rights reserved

VOLUME 63 / ISSUE 3 MARCH 2026

Contents

22
AI Hunts for New Physics
Machine learning may spot patterns that humans never imagined.
By Matthew Hutson

Particle Physics Seeks a Cancer Cure 32
A split second of ultrahigh-power radiation kills tumors.
By Tom Clynes

The Anti-Induction Suit 40
High-tech clothing protects power-line workers.
By Peter Fairley

In 1844, Chess Was Already Online 42
From telegraph to esport, chess play has continued to evolve.
By Greg Uyeno

THE INSTITUTE 62
From Stage and Screen to AI Consulting
Gokul Pandy is a leading man at Accenture.

EDITOR'S NOTE 2
AI is becoming a scientific instrument.

NEWS 5
3D Chip Cooling
Stretchable OLEDs
Power Beaming From Planes

ON THE COVER:
Illustration by Peter Crowther

THE DATA 14
Putting a data center in orbit would be expensive.

HANDS ON 17
Re-creating Jimi Hendrix's system engineering.

CAREERS 19
This ex-physicist keeps major networks secure.

5 QUESTIONS 21
Laslo Hunhold invents a new number format.

PAST FORWARD 68
An Early EV Enthusiast

CERN

EDITOR'S NOTE

BY ELIZA STRICKLAND

Letting Machines Decide What Matters

AI systems built into particle detectors now shape what physicists study

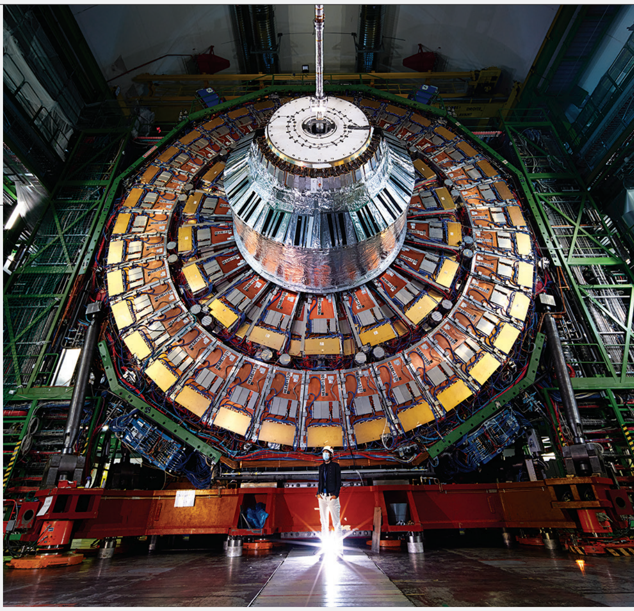

This CERN particle detector uses AI to decide what data gets saved.

In the time it takes you to read this sentence, the Large Hadron Collider (LHC) will have smashed billions of particles together. In all likelihood, it will have found exactly what it found yesterday: more evidence to support the Standard Model of particle physics.

For the engineers who built the 27-kilometer-long ring, this consistency is a triumph. But for theoretical physicists, it has been rather frustrating. As Matthew Hutson reports in “AI Hunts for the Next Big Thing in Physics” [p. 22], the field is currently gripped by a quiet crisis. In an email discussing his reporting, Hutson explains that the Standard Model, which describes the known elementary particles and forces, is not a complete picture. “So theorists have proposed new ideas, and experimentalists have built giant facilities to test them, but despite the gobs of data, there have been no big breakthroughs,” Hutson says. “There are key components of reality we’re completely missing.”

> “There are key components of reality we’re completely missing.”
>
> —Matthew Hutson

That’s why researchers are turning artificial intelligence loose on particle physics. They aren’t simply asking AI to comb through accelerator data to confirm existing theories, Hutson explains. They’re asking AI to point the way toward theories that they’ve never imagined. “Instead of looking to support theories that humans have generated,” he says, “unsupervised AI can highlight anything out of the ordinary, expanding our reach into unknown unknowns.” By asking AI to flag anomalies in the data, researchers hope to find their way to “new physics” that extends the Standard Model.

On the surface, this article might sound like another “AI for *X*” story. As *IEEE Spectrum*’s AI editor, I get a steady stream of pitches for such stories: AI for drug discovery, AI for farming, AI for wildlife tracking. Often what that really means is faster data processing or automation around the edges. Useful, sure, but incremental.

What struck me in Hutson’s reporting is that this effort feels different. Instead of analyzing experimental data after the fact, the AI essentially becomes part of the instrument, scanning for subtle patterns and deciding in real time what’s interesting. At the LHC, detectors record 40 million collisions per second. There’s simply no way to preserve all that data, so engineers have always had to build filters to decide which events get saved for analysis and which are discarded; nearly everything is thrown away. Now those split-second decisions are increasingly handed to machine learning systems running on field-programmable gate arrays (FPGAs) connected to the detectors. The code must run on the chip’s limited logic and memory, and compressing a neural network into that hardware isn’t easy. Hutson describes one theorist pleading with an engineer, “Which of my algorithms fits on your bloody FPGA?”

This moment is part of a much older pattern. As Hutson writes in the article, new instruments have opened doors to the unexpected throughout the history of science. Galileo’s telescope revealed moons circling Jupiter. Early microscopes exposed entire worlds of “animalcules” swimming around. These better tools didn’t just answer existing questions; they made it possible to ask new ones.

If there’s a crisis in particle physics, in other words, it may not just be about missing particles. It’s about how to look beyond the limits of the human imagination. Hutson’s story suggests that AI might not solve the mysteries of the universe outright, but it could change how we search for answers. ■

LEFT: NICOLE MILLMAN; RIGHT: CERN

A New Prize Category

THE SHAW PRIZE IN COMPUTER SCIENCE

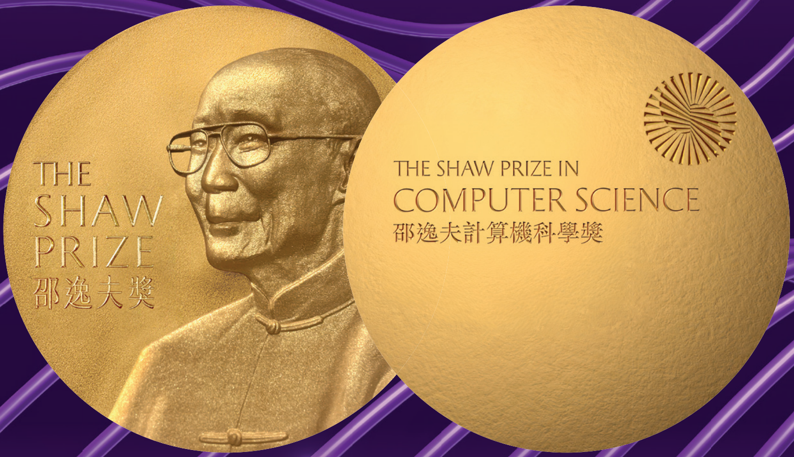
THE SHAW PRIZE
邵逸夫獎
THE SHAW PRIZE IN
COMPUTER SCIENCE
邵逸夫計算機科學獎

Nomination starts in September 2026

First winner(s) will be announced in spring 2027

SELECTION COMMITTEE

PROFESSOR
JENNIFER T CHAYES

Chair

PROFESSOR
TONY CHAN

PROFESSOR
JACK DONGARRA

PROFESSOR
JOHN L HENNESSY

PROFESSOR
YANN LECUN

PROFESSOR
HARRY SHUM

PROFESSOR
JOSEPH SIFAKIS

PROFESSOR
JEANNETTE M WING

shawprize.org

CONTRIBUTORS

TOM CLYNES

In this issue, Clynes, an *IEEE Spectrum* contributing editor, looks into a promising new cancer treatment [p. 32]. FLASH therapy replaces months of conventional cancer treatment with a split-second blast of ultrahigh-power radiation—delivered by particle accelerators engineered for extreme beam precision. Clynes visited labs in France, Germany, and Switzerland where the technology is being developed. "I was struck by the quiet confidence within the research community that this approach—which runs counter to established radiobiology doctrine—is likely to be the dominant way to treat tumors within a decade," Clynes says.

MATTHEW HUTSON

At CERN's Large Hadron Collider, nearly all the data gets thrown away. But what if some of that discarded data contains clues to "new physics" that push beyond today's theories? In this issue, Hutson, a New York–based science journalist, describes how particle physicists are using AI to search for those signals [p. 22]. By integrating AI into systems that decide what data to keep, they're asking algorithms to surface patterns no one thought to look for. "The idea of asking AI to decide what's interesting fascinated me," Hutson says.

ANDREW MOSEMAN

Moseman is a science and tech journalist who primarily covers energy tech and EVs. In this issue, he writes about Overview Energy, a startup that successfully beamed power from an airplane to the ground [p. 9]. He says he was surprised at the company's practical workaround to existing challenges, by using frequencies that aren't in "the crowded 'beachfront' of the wireless spectrum, and being able to use garden-variety solar panels to receive it. That's pretty cool."

GREG UYENO

Uyeno is a freelance science journalist based in New York City. In this issue, he writes about the history of telegraph chess [p. 42], starting with the first such match in 1844. Uyeno traces how chess has continued to migrate to other electronic platforms, most recently as an esport. "Chess players have consistently adopted telecommunications as a convenience, a novelty, and a publicity boon," Uyeno says. "And in turn, chess has helped to test, promote, or legitimize nascent technologies."

IEEE Spectrum

EDITOR IN CHIEF Harry Goldstein, h.goldstein@ieee.org
EXECUTIVE EDITOR Jean Kumagai, j.kumagai@ieee.org
MANAGING EDITOR Elizabeth A. Bretz, e.bretz@ieee.org
CREATIVE DIRECTOR Mark Montgomery, m.montgomery@ieee.org
DIRECTOR OF DIGITAL INNOVATION Erico Guizzo, e.guizzo@ieee.org
EDITORIAL DIRECTOR, CONTENT DEVELOPMENT
Glenn Zorpette, g.zorpette@ieee.org

SENIOR EDITORS
Evan Ackerman, evan.ackerman@ieee.org
Lucas Laursen, l.laursen@ieee.org
Samuel K. Moore, s.k.moore@ieee.org
Eliza Strickland, e.strickland@ieee.org

NEWS MANAGER Michael Koziol, m.koziol@ieee.org

SENIOR ASSOCIATE EDITOR Emily Waltz, e.waltz@ieee.org

ASSOCIATE EDITORS
Dina Genkina, d.genkina@ieee.org
Willie D. Jones, w.jones@ieee.org

ASSISTANT EDITOR Gwendolyn Rak, g.rak@ieee.org

SPECIAL PROJECTS DIRECTOR
Stephen Cass, cass.s@ieee.org

SENIOR ASSOCIATE EDITOR
Margo Anderson, m.k.anderson@ieee.org

PRODUCT AND AUDIENCE DEVELOPMENT DIRECTOR
Tim Warder, t.warder@ieee.org
PRODUCT MANAGER Priyanka Shitole, p.shitole@ieee.org
AUDIENCE DEVELOPMENT MANAGER
Laura Bridgeman, l.bridgeman@ieee.org
COMMUNITY MANAGER
Kohava Mendelsohn, k.mendelsohn@ieee.org
DIGITAL PRODUCT DESIGNER Erik Vrielink, e.vrielink@ieee.org

ART & PRODUCTION
ART DIRECTOR Brandon Palacio, b.palacio@ieee.org
PHOTOGRAPHY DIRECTOR Randi Klett, randi.klett@ieee.org
ASSISTANT PHOTO EDITOR Nicole Millman, n.millman@ieee.org
PRINT PRODUCTION SPECIALIST
Sylvana Meneses, s.meneses@ieee.org
DIGITAL PRODUCTION MANAGER
Michael Spector, m.spector@ieee.org

SENIOR COPY EDITOR Joseph N. Levine, j.levine@ieee.org
COPY EDITOR Michele Kogon, m.kogon@ieee.org
RESEARCH COPY EDITOR Cheryl Della Pietra, c.dellapietra@ieee.org
ADMINISTRATOR Jose Zambrano, jose.zambrano@ieee.org
CONTRIBUTING EDITORS Rina Diane Caballar, Robert N. Charette, Charles Q. Choi, Tom Clynes, Peter Fairley, Edd Gent, W. Wayt Gibbs, Mark Harris, Allison Marsh, Julianne Pepitone, Ned Potter, Matthew S. Smith, Lawrence Ulrich

THE INSTITUTE
EDITOR IN CHIEF Kathy Pretz, k.pretz@ieee.org
ASSOCIATE EDITOR Joanna Goodrich, j.goodrich@ieee.org

SUBSCRIPTION SERVICES
IEEE Members member-services@ieee.org
Non-members +1 866 363 2304

DIRECTOR, PERIODICALS PRODUCTION SERVICES Peter Tuohy
ADVERTISING PRODUCTION MANAGER
Felicia Spagnoli, f.spagnoli@ieee.org
ADVERTISING PRODUCTION +1 732 562 6334

EDITORIAL ADVISORY BOARD, *IEEE SPECTRUM*
Harry Goldstein, *Chair*; Muhammad Alaa, Karanveer Anand, Ella M. Atkins, Sangyeun Cho, Hugh Durrant-Whyte, Benjamin Gross, Carlos Gutierrez, Lawrence O. Hall, Benjamin Kroposki, Somdeb Majumdar, Khaled Mokhtar, Ramune Nagisetty, Paul Nielsen, Ketra Schmidt, Gaurav Sharma, Fan Shi, Mini S. Thomas, Andrej Zdravkovic, Lidong Zhou

EDITORIAL ADVISORY BOARD, *THE INSTITUTE*
Kathy Pretz, *Chair*; Stamatis Dragoumanos, Madeleine Glick, Harry Goldstein, Lawrence O. Hall, Francesca Iacopi, Sai Prashanth Mallellu, Shashi Raj Pandey, John Purvis, Chenyang Xu

EDITORIAL CORRESPONDENCE
IEEE Spectrum, 3 Park Ave., 17th Floor,
New York, NY 10016-5997 Tel: +1 212 419 7555

ADVERTISING INQUIRIES Wiley
Joerg Wuellner +49 152 2258 4972, jwuellner@wiley.com
Chris Darch +1 206 867 4165, cdarch@wiley.com

REPRINT SALES +1 212 221 9595, ext. 319

REPRINT PERMISSION/ LIBRARIES Articles may be photocopied for private use of patrons. A per-copy fee must be paid to the Copyright Clearance Center, 29 Congress St., Salem, MA 01970. For other copying or republication, contact Managing Editor, *IEEE Spectrum.*

COPYRIGHTS AND TRADEMARKS *IEEE Spectrum* is a registered trademark owned by The Institute of Electrical and Electronics Engineers Inc. Responsibility for the substance of articles rests upon the authors, not IEEE, its organizational units, or its members. Articles do not represent official positions of IEEE. Readers may post comments online; comments may be excerpted for publication. IEEE reserves the right to reject any advertising.

IEEE

IEEE BOARD OF DIRECTORS
PRESIDENT&CEO Mary Ellen Randall, president@ieee.org
+1 732 562 3928 Fax: +1 732 981 9515
PRESIDENT-ELECT Jill I. Gostin
PAST PRESIDENT Kathleen A. Kramer
TREASURER Gerardo Barbosa
SECRETARY Jenifer Castillo Rodriguez

VICE PRESIDENTS
Timothy P. Kurzweg, *Educational Activities;* Andrew D. Lowery, *Member & Geographic Activities;* W. Clem Karl, *Publication Services & Products;* F.D. Tan, *Technical Activities;* Gary R. Hoffman, *President, Standards Association;* Barry C. Tilton, *President, IEEE-USA*

DIVISION DIRECTORS
Fernando J. Guarin (I); H. Alan Mantooth (II); Sarah Kate Wilson (III); Charles M. Jackson (IV); Leila De Floriani (V); Russell D. Meier (VI); Noel M. Schulz (VII); Cecilia Metra (VIII); Cristina Olaverri-Monreal (IX); Ravinder S. Dahiya (X)

REGION DIRECTORS
Charles P. Rubenstein (1); Felicia Reinhart (2); Sonya L. Dillard (3); Constance A. Kelly (4); Christopher B. Sanderson (5); Joseph C. Wei (6); Wahab Almuhtadi (7); Dusanka Boskovic (8); Jose-Ignacio Castillo-Velazquez (9); Takako Hashimoto (10)

IEEE STAFF
EXECUTIVE DIRECTOR & COO Sophia A. Muirhead
+1 732 562 5400, s.muirhead@ieee.org
CHIEF FINANCIAL OFFICER Kelly Armstrong
+1 732 562 6011, k.armstrong@ieee.org
MANAGING DIRECTOR,TECHNICAL ACTIVITIES Mojdeh Bahar
+1 732 562 3850, m.bahar@ieee.org
GENERAL COUNSEL & CHIEF COMPLIANCE OFFICER
Ahsaki Benion +1 212 705 8975, a.benion@ieee.org
MANAGING DIRECTOR, PUBLICATIONS Steven Heffner
+1 212 705 8958, s.heffner@ieee.org
CHIEF GOVERNANCE OFFICER Donna Hourican
+1 732 562 6330, d.hourican@ieee.org
MANAGING DIRECTOR, CONFERENCES, EVENTS AND EXPERIENCES Marie Hunter
+ 1 732 465 5889, m.m.hunter@ieee.org
MANAGING DIRECTOR, MEMBER & GEOGRAPHIC ACTIVITIES
Cecelia Jankowski +1 732 562 5504, c.jankowski@ieee.org
MANAGING DIRECTOR, EDUCATIONAL ACTIVITIES
Jamie Moesch +1 732 562 5514, j.moesch@ieee.org
CHIEF MARKETING OFFICER Jayne O'Brien
+1 732 981 3433, jayne.obrien@ieee.org
MANAGING DIRECTOR, IEEE STANDARDS Alpesh Shah
+1 732 465 6467, alpesh.shah@ieee.org
CHIEF INFORMATION AND DIGITAL OFFICER Jeff Strohschein
+1 732 562 5332, j.strohschein@ieee.org
CHIEF OF STAFF Allison Taylor
+1 732 562 6068, allison.taylor@ieee.org
CHIEF HUMAN RESOURCES OFFICER Cheri N. Collins Wideman
+1 732 562 5335, c.wideman@ieee.org
MANAGING DIRECTOR, IEEE-USA Russell T. Harrison
+1 202 530 8326, r.t.harrison@ieee.org

IEEE PUBLICATION SERVICES & PRODUCTS BOARD
W. Clem Karl, *Chair*; John T. Barr, Nazanin Bassiri-Gharb, Sergio Benedetto, Eddie Custovic, Stephen Dukes, Alessandro Foi, Lawrence O. Hall, Charles Hansen, Warunika Hippola, James Irvine, Michael Jensen, Athanasios Kakarountas, James Keller, Imran Mehdi, Khaled Mokhtar, Paolo Montuschi, Zorana Popovic, Mehrdad Saif, Gianluca Setti, Gaurav Sharma, Sarah Spurgeon, Jaideep Vaidya, Maria Elena Valcher, Patrick Wheeler, Steve Yurkovich

IEEE OPERATIONS CENTER
445 Hoes Lane, Box 1331, Piscataway, NJ 08854-1331 U.S.A.
Tel: +1 732 981 0060 Fax: +1 732 981 1721

IEEE SPECTRUM (ISSN 0018-9235) is published monthly by The Institute of Electrical and Electronics Engineers, Inc. All rights reserved. © 2026 by The Institute of Electrical and Electronics Engineers, Inc., 3 Park Avenue, New York, NY 10016-5997, U.S.A. Volume No. 63, Issue No. 3. The editorial content of *IEEE Spectrum* magazine does not represent official positions of the IEEE or its organizational units. Canadian Post International Publications Mail (Canadian Distribution) Sales Agreement No. 40013087. Return undeliverable Canadian addresses to: Circulation Department, *IEEE Spectrum*, Box 1051, Fort Erie, ON L2A 6C7. Cable address: ITRIPLEE. Fax: +1 212 419 7570. INTERNET: spectrum@ieee.org. ANNUAL SUBSCRIPTIONS: IEEE Members: $21.40 included in dues. Libraries/institutions: $399. POSTMASTER: Please send address changes to *IEEE Spectrum*, % Coding Department, IEEE Service Center, 445 Hoes Lane, Box 1331, Piscataway, NJ 08855. Periodicals postage paid at New York, NY, and additional mailing offices. Canadian GST #125634188. Printed at 3401 N. Heartland Dr., Liberty, MO 64068, U.S.A. Printed at Spenta Multimedia Pvt. Ltd., Plot 15, 16 & 21/1, Village Chikhloli, Morivali, MIDC, Ambernath (West), Dist. Thane. *IEEE Spectrum* circulation is audited by BPA Worldwide. *IEEE Spectrum* is a member of the Association of Business Information & Media Companies, the Association of Magazine Media, and Association Media & Publishing. IEEE prohibits discrimination, harassment, and bullying. For more information, visit https://www.ieee.org/about/corporate/governance/p9-26.html.

CONTRIBUTORS

TOM CLYNES

In this issue, Clynes, an *IEEE Spectrum* contributing editor, looks into a promising new cancer treatment [p. 32]. FLASH therapy replaces months of conventional cancer treatment with a split-second blast of ultrahigh-power radiation—delivered by particle accelerators engineered for extreme beam precision. Clynes visited labs in France, Germany, and Switzerland where the technology is being developed. "I was struck by the quiet confidence within the research community that this approach—which runs counter to established radiobiology doctrine—is likely to be the dominant way to treat tumors within a decade," Clynes says.

MATTHEW HUTSON

At CERN's Large Hadron Collider, nearly all the data gets thrown away. But what if some of that discarded data contains clues to "new physics" that push beyond today's theories? In this issue, Hutson, a New York–based science journalist, describes how particle physicists are using AI to search for those signals [p. 22]. By integrating AI into systems that decide what data to keep, they're asking algorithms to surface patterns no one thought to look for. "The idea of asking AI to decide what's interesting fascinated me," Hutson says.

ANDREW MOSEMAN

Moseman is a science and tech journalist who primarily covers energy tech and EVs. In this issue, he writes about Overview Energy, a startup that successfully beamed power from an airplane to the ground [p. 9]. He says he was surprised at the company's practical workaround to existing challenges, by using frequencies that aren't in "the crowded 'beachfront' of the wireless spectrum, and being able to use garden-variety solar panels to receive it. That's pretty cool."

GREG UYENO

Uyeno is a freelance science journalist based in New York City. In this issue, he writes about the history of telegraph chess [p. 42], starting with the first such match in 1844. Uyeno traces how chess has continued to migrate to other electronic platforms, most recently as an esport. "Chess players have consistently adopted telecommunications as a convenience, a novelty, and a publicity boon," Uyeno says. "And in turn, chess has helped to test, promote, or legitimize nascent technologies."

IEEE Spectrum

EDITOR IN CHIEF Harry Goldstein, h.goldstein@ieee.org
EXECUTIVE EDITOR Jean Kumagai, j.kumagai@ieee.org
MANAGING EDITOR Elizabeth A. Bretz, e.bretz@ieee.org
CREATIVE DIRECTOR Mark Montgomery, m.montgomery@ieee.org
DIRECTOR OF DIGITAL INNOVATION Erico Guizzo, e.guizzo@ieee.org
EDITORIAL DIRECTOR, CONTENT DEVELOPMENT
Glenn Zorpette, g.zorpette@ieee.org

SENIOR EDITORS
Evan Ackerman, evan.ackerman@ieee.org
Lucas Laursen, l.laursen@ieee.org
Samuel K. Moore, s.k.moore@ieee.org
Eliza Strickland, e.strickland@ieee.org

NEWS MANAGER Michael Koziol, m.koziol@ieee.org

SENIOR ASSOCIATE EDITOR Emily Waltz, e.waltz@ieee.org

ASSOCIATE EDITORS
Dina Genkina, d.genkina@ieee.org
Willie D. Jones, w.jones@ieee.org

ASSISTANT EDITOR Gwendolyn Rak, g.rak@ieee.org

SPECIAL PROJECTS DIRECTOR
Stephen Cass, cass.s@ieee.org

SENIOR ASSOCIATE EDITOR
Margo Anderson, m.k.anderson@ieee.org

PRODUCT AND AUDIENCE DEVELOPMENT DIRECTOR
Tim Warder, t.warder@ieee.org
PRODUCT MANAGER Priyanka Shitole, p.shitole@ieee.org
AUDIENCE DEVELOPMENT MANAGER
Laura Bridgeman, l.bridgeman@ieee.org
COMMUNITY MANAGER
Kohava Mendelsohn, k.mendelsohn@ieee.org
DIGITAL PRODUCT DESIGNER Erik Vrielink, e.vrielink@ieee.org

ART & PRODUCTION
ART DIRECTOR Brandon Palacio, b.palacio@ieee.org
PHOTOGRAPHY DIRECTOR Randi Klett, randi.klett@ieee.org
ASSISTANT PHOTO EDITOR Nicole Millman, n.millman@ieee.org
PRINT PRODUCTION SPECIALIST
Sylvana Meneses, s.meneses@ieee.org
DIGITAL PRODUCTION MANAGER
Michael Spector, m.spector@ieee.org

SENIOR COPY EDITOR Joseph N. Levine, j.levine@ieee.org
COPY EDITOR Michele Kogon, m.kogon@ieee.org
RESEARCH COPY EDITOR Cheryl Della Pietra, c.dellapietra@ieee.org
ADMINISTRATOR Jose Zambrano, jose.zambrano@ieee.org
CONTRIBUTING EDITORS Rina Diane Caballar, Robert N. Charette, Charles Q. Choi, Tom Clynes, Peter Fairley, Edd Gent, W. Wayt Gibbs, Mark Harris, Allison Marsh, Julianne Pepitone, Ned Potter, Matthew S. Smith, Lawrence Ulrich

THE INSTITUTE
EDITOR IN CHIEF Kathy Pretz, k.pretz@ieee.org
ASSOCIATE EDITOR Joanna Goodrich, j.goodrich@ieee.org

SUBSCRIPTION SERVICES
IEEE Members member-services@ieee.org
Non-members +1 866 363 2304

DIRECTOR, PERIODICALS PRODUCTION SERVICES Peter Tuohy
ADVERTISING PRODUCTION MANAGER
Felicia Spagnoli, f.spagnoli@ieee.org
ADVERTISING PRODUCTION +1 732 562 6334

EDITORIAL ADVISORY BOARD, *IEEE SPECTRUM*
Harry Goldstein, *Chair*; Muhammad Alaa, Karanveer Anand, Ella M. Atkins, Sangyeun Cho, Hugh Durrant-Whyte, Benjamin Gross, Carlos Gutierrez, Lawrence O. Hall, Benjamin Kroposki, Somdeb Majumdar, Khaled Mokhtar, Ramune Nagisetty, Paul Nielsen, Ketra Schmidt, Gaurav Sharma, Fan Shi, Mini S. Thomas, Andrej Zdravkovic, Lidong Zhou

EDITORIAL ADVISORY BOARD, *THE INSTITUTE*
Kathy Pretz, *Chair*; Stamatis Dragoumanos, Madeleine Glick, Harry Goldstein, Lawrence O. Hall, Francesca Iacopi, Sai Prashanth Mallellu, Shashi Raj Pandey, John Purvis, Chenyang Xu

EDITORIAL CORRESPONDENCE
IEEE Spectrum, 3 Park Ave., 17th Floor,
New York, NY 10016-5997 Tel: +1 212 419 7555

ADVERTISING INQUIRIES Wiley
Joerg Wuellner +49 152 2258 4972, jwuellner@wiley.com
Chris Darch +1 206 867 4165, cdarch@wiley.com

REPRINT SALES +1 212 221 9595, ext. 319

REPRINT PERMISSION/ LIBRARIES Articles may be photocopied for private use of patrons. A per-copy fee must be paid to the Copyright Clearance Center, 29 Congress St., Salem, MA 01970. For other copying or republication, contact Managing Editor, *IEEE Spectrum.*

COPYRIGHTS AND TRADEMARKS *IEEE Spectrum* is a registered trademark owned by The Institute of Electrical and Electronics Engineers Inc. Responsibility for the substance of articles rests upon the authors, not IEEE, its organizational units, or its members. Articles do not represent official positions of IEEE. Readers may post comments online; comments may be excerpted for publication. IEEE reserves the right to reject any advertising.

IEEE

IEEE BOARD OF DIRECTORS
PRESIDENT&CEO Mary Ellen Randall, president@ieee.org
+1 732 562 3928 Fax: +1 732 981 9515
PRESIDENT-ELECT Jill I. Gostin
PAST PRESIDENT Kathleen A. Kramer
TREASURER Gerardo Barbosa
SECRETARY Jenifer Castillo Rodriguez

VICE PRESIDENTS
Timothy P. Kurzweg, *Educational Activities;* Andrew D. Lowery, *Member & Geographic Activities;* W. Clem Karl, *Publication Services & Products;* F.D. Tan, *Technical Activities;* Gary R. Hoffman, *President, Standards Association;* Barry C. Tilton, *President, IEEE-USA*

DIVISION DIRECTORS
Fernando J. Guarin (I); H. Alan Mantooth (II); Sarah Kate Wilson (III); Charles M. Jackson (IV); Leila De Floriani (V); Russell D. Meier (VI); Noel M. Schulz (VII); Cecilia Metra (VIII); Cristina Olaverri-Monreal (IX); Ravinder S. Dahiya (X)

REGION DIRECTORS
Charles P. Rubenstein (1); Felicia Reinhart (2); Sonya L. Dillard (3); Constance A. Kelly (4); Christopher B. Sanderson (5); Joseph C. Wei (6); Wahab Almuhtadi (7); Dusanka Boskovic (8); Jose-Ignacio Castillo-Velazquez (9); Takako Hashimoto (10)

IEEE STAFF
EXECUTIVE DIRECTOR & COO Sophia A. Muirhead
+1 732 562 5400, s.muirhead@ieee.org
CHIEF FINANCIAL OFFICER Kelly Armstrong
+1 732 562 6011, k.armstrong@ieee.org
MANAGING DIRECTOR,TECHNICAL ACTIVITIES Mojdeh Bahar
+1 732 562 3850, m.bahar@ieee.org
GENERAL COUNSEL & CHIEF COMPLIANCE OFFICER
Ahsaki Benion +1 212 705 8975, a.benion@ieee.org
MANAGING DIRECTOR, PUBLICATIONS Steven Heffner
+1 212 705 8958, s.heffner@ieee.org
CHIEF GOVERNANCE OFFICER Donna Hourican
+1 732 562 6330, d.hourican@ieee.org
MANAGING DIRECTOR, CONFERENCES, EVENTS AND EXPERIENCES Marie Hunter
+ 1 732 465 5889, m.m.hunter@ieee.org
MANAGING DIRECTOR, MEMBER & GEOGRAPHIC ACTIVITIES
Cecelia Jankowski +1 732 562 5504, c.jankowski@ieee.org
MANAGING DIRECTOR, EDUCATIONAL ACTIVITIES
Jamie Moesch +1 732 562 5514, j.moesch@ieee.org
CHIEF MARKETING OFFICER Jayne O'Brien
+1 732 981 3433, jayne.obrien@ieee.org
MANAGING DIRECTOR, IEEE STANDARDS Alpesh Shah
+1 732 465 6467, alpesh.shah@ieee.org
CHIEF INFORMATION AND DIGITAL OFFICER Jeff Strohschein
+1 732 562 5332, j.strohschein@ieee.org
CHIEF OF STAFF Allison Taylor
+1 732 562 6068, allison.taylor@ieee.org
CHIEF HUMAN RESOURCES OFFICER Cheri N. Collins Wideman
+1 732 562 5335, c.wideman@ieee.org
MANAGING DIRECTOR, IEEE-USA Russell T. Harrison
+1 202 530 8326, r.t.harrison@ieee.org

IEEE PUBLICATION SERVICES & PRODUCTS BOARD
W. Clem Karl, *Chair*; John T. Barr, Nazanin Bassiri-Gharb, Sergio Benedetto, Eddie Custovic, Stephen Dukes, Alessandro Foi, Lawrence O. Hall, Charles Hansen, Warunika Hippola, James Irvine, Michael Jensen, Athanasios Kakarountas, James Keller, Imran Mehdi, Khaled Mokhtar, Paolo Montuschi, Zorana Popovic, Mehrdad Saif, Gianluca Setti, Gaurav Sharma, Sarah Spurgeon, Jaideep Vaidya, Maria Elena Valcher, Patrick Wheeler, Steve Yurkovich

IEEE OPERATIONS CENTER
445 Hoes Lane, Box 1331, Piscataway, NJ 08854-1331 U.S.A.
Tel: +1 732 981 0060 Fax: +1 732 981 1721

IEEE SPECTRUM (ISSN 0018-9235) is published monthly by The Institute of Electrical and Electronics Engineers, Inc. All rights reserved. © 2026 by The Institute of Electrical and Electronics Engineers, Inc., 3 Park Avenue, New York, NY 10016-5997, U.S.A. Volume No. 63, Issue No. 3. The editorial content of *IEEE Spectrum* magazine does not represent official positions of the IEEE or its organizational units. Canadian Post International Publications Mail (Canadian Distribution) Sales Agreement No. 40013087. Return undeliverable Canadian addresses to: Circulation Department, *IEEE Spectrum*, Box 1051, Fort Erie, ON L2A 6C7. Cable address: ITRIPLEE. Fax: +1 212 419 7570. INTERNET: spectrum@ieee.org. ANNUAL SUBSCRIPTIONS: IEEE Members: $21.40 included in dues. Libraries/institutions: $399. POSTMASTER: Please send address changes to *IEEE Spectrum*, % Coding Department, IEEE Service Center, 445 Hoes Lane, Box 1331, Piscataway, NJ 08855. Periodicals postage paid at New York, NY, and additional mailing offices. Canadian GST #125634188. Printed at 3401 N. Heartland Dr., Liberty, MO 64068, U.S.A. Printed at Spenta Multimedia Pvt. Ltd., Plot 15, 16 & 21/1, Village Chikhloli, Morivali, MIDC, Ambernath (West), Dist. Thane. *IEEE Spectrum* circulation is audited by BPA Worldwide. *IEEE Spectrum* is a member of the Association of Business Information & Media Companies, the Association of Magazine Media, and Association Media & Publishing. IEEE prohibits discrimination, harassment, and bullying. For more information, visit https://www.ieee.org/about/corporate/governance/p9-26.html.

News

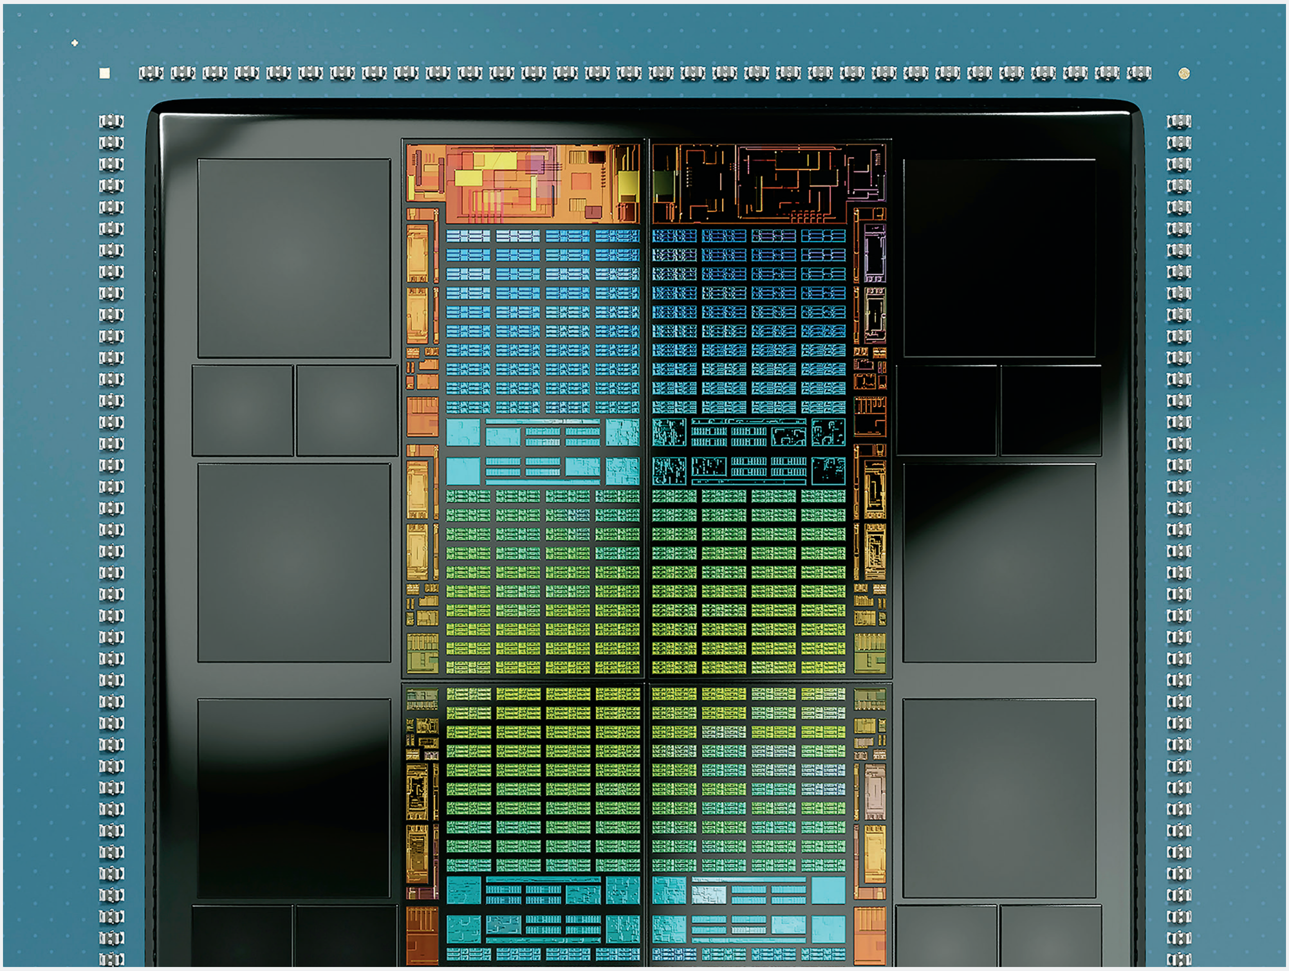

High-bandwidth memory [black squares] is currently attached to the sides of a GPU's logic die, as in this Instinct MI300X GPU from AMD.

AMD

SEMICONDUCTORS

3D Chip Integration Risks Cooking Future GPUs › Imec has a multistep plan to keep things cool

BY SAMUEL K. MOORE

Peek inside the package of Advanced Micro Devices' or Nvidia's most advanced AI products and you'll find a familiar arrangement: The GPU is flanked on two sides by high-bandwidth memory (HBM), the most advanced memory chips available. These DRAM chips are placed as close as possible to the computing chips they serve in order to cut down on the biggest bottleneck in AI computing—the energy and delay in getting billions of bits per second from memory into logic. But what if you could bring computing and memory even closer together by stacking the HBM on top of the GPU?

Imec recently explored this scenario using advanced thermal simulations, and the answer—delivered in December at the 2025 IEEE International Electron Devices Meeting (IEDM)—was a bit grim. 3D stacking doubles the temperature inside the GPU, rendering it inoperable. But the team, led by James Myers, a program director at Imec, didn't just give up. They identified several engineering optimizations that ultimately could whittle down the temperature difference from stacking chips to nearly zero.

Imec started with a thermal simulation of a GPU and four HBM dies as you'd find them today, inside what's called a 2.5D package. In this configuration, the GPU and the four HBMs sit on a substrate called an interposer, with minimal distance between them to cut down on delays. The two types of chips are linked by thousands of micrometer-scale copper interconnects built into the interposer's surface. In this configuration, the model GPU consumes 414 watts and reaches a peak temperature of just under 70 °C—typical for a processor. The memory chips consume an additional 40 W or so and heat up slightly less than the GPU. The heat is removed from the top of the package by the kind of liquid cooling that's become common in new AI data centers.

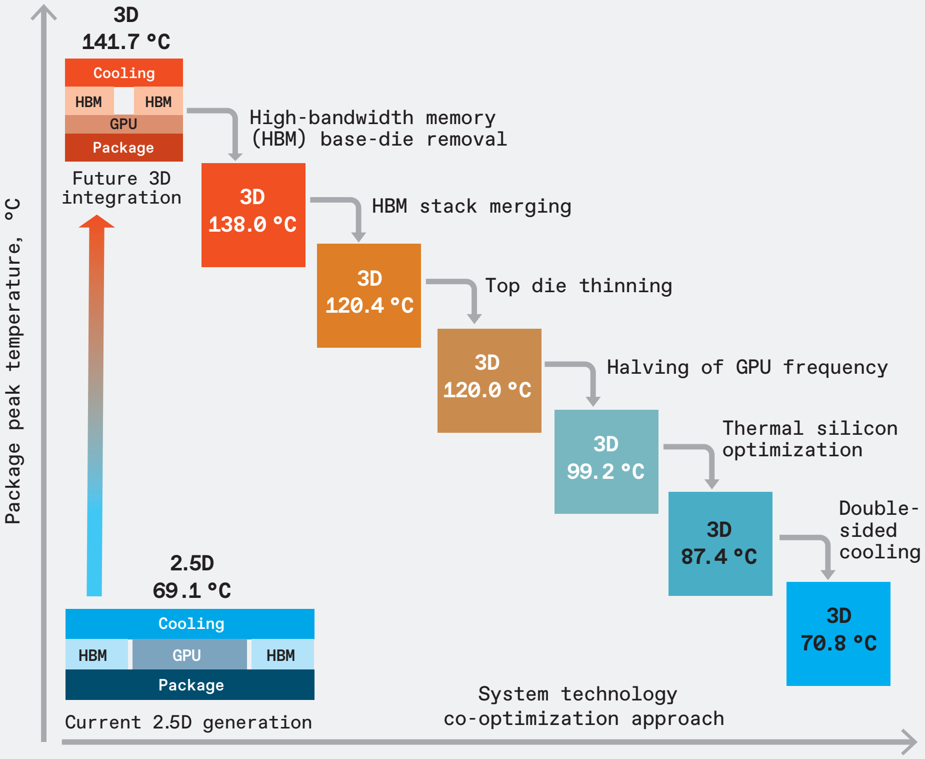

Imec has developed a multistep process to reduce the heat generated by 3D chip stacking so that it is comparable to existing 2.5D solutions.

"While this approach is currently used, it does not scale well for the future—especially as it blocks two sides of the GPU, limiting future GPU-to-GPU connections inside the package," Yukai Chen, a senior researcher at Imec, told engineers at IEDM. By contrast, he says, "the 3D approach leads to higher bandwidth, lower latency.…The most important improvement is the package footprint."

Unfortunately, as Chen and his colleagues found, the most straightforward version of stacking—simply putting the HBM chips on top of the GPU and adding a block of blank silicon to fill in a gap at the center—cuts off the path heat usually uses to escape from the GPU. This shoots up temperatures in the processor to a scorching 140 °C—well past a typical GPU's 80 °C limit.

The Imec team set about trying a number of technology and system optimizations aimed at lowering the temperature. The first thing they tried was throwing out a layer of silicon that was now redundant. To understand why, you have to first get a grip on what HBM really is.

This form of memory is a stack of as many as 12 high-density DRAM dies. Each has been thinned down to tens of micrometers and is shot through with vertical connections. These thinned dies are stacked one atop another and connected by tiny balls of solder, and this stack of memory is vertically connected to another piece of silicon, called the base die. The base die is a logic chip designed to multiplex the data—in other words, to pack it into the limited number of wires that can fit across the millimeter-scale gap to the GPU.

But with the HBM now on top of the GPU, there's no need for such a data pump. Bits can flow directly into the processor regardless of how many wires happen to fit along the side of the chip. This change also means moving the memory control circuits from the base die into the GPU, and therefore changing the processor's floor plan, says Myers. But there should be ample room, he suggests, because the GPU will no longer need the circuits used to demultiplex incoming memory data.

Slowing the GPU's clock by 50 percent leads to a performance win, while cooling everything down by more than 20 °C.

AMD

Cutting out this middleman of memory cooled things down by only a little less than 4 °C. But it should massively boost the bandwidth between the memory and the processor, which is important for another optimization the team tried—slowing down the GPU.

That might seem contrary to the whole purpose of better AI computing, but in this case, it's an advantage. Large language models are considered "memory-bound" problems—meaning that memory bandwidth is the main limiting factor to their speed. But Myers's team estimated that 3D-stacking HBM on the GPU would boost bandwidth fourfold. With that added headroom, even slowing the GPU's clock by 50 percent still leads to a performance win, while cooling everything down by more than 20 °C. In practice, the processor might not need to be slowed down quite that much: Slowing the clock frequency by just 30 percent led to a GPU that was only 1.7 °C warmer, Myers says.

Another big drop in temperature came from making the HBM stack and the area around it more conductive to heat. That included merging the four stacks into two wider stacks, thereby eliminating a heat-trapping region; thinning out the top die of the stack, which is usually thicker; and filling in more of the space around the HBM with blank pieces of silicon to conduct more heat.

With all of that, the stack now operated at about 88 °C. One final optimization brought things back to near 70 °C. Generally, some 95 percent of a chip's heat is removed from the top of the package, where in this case water carries the heat away. But adding similar cooling to the underside as well drove the stacked chips down a final 17 °C.

Although the research presented at IEDM shows it might be possible, HBM-on-GPU isn't necessarily the right way to go, Myers says. "We are simulating other system configurations to help build confidence that this is or isn't the best choice," he says. "GPU-on-HBM is of interest to some in industry," because it puts the GPU closer to the cooling. But it would likely be a more complex design, because the GPU's power and data would have to flow vertically through the HBM to reach it. ■

ARTIFICIAL INTELLIGENCE

AI Helps Scientists but Hurts Science

› New analysis suggests AI tools flatten discovery

BY ELIE DOLGIN

Artificial intelligence is turning scientists into publishing machines. It's also quietly funneling them into the same crowded corners of research.

That's what a group of researchers in China and the United States concluded after analyzing more than 40 million academic papers. They found that scientists who use AI tools in their research publish more papers, accumulate more citations, and reach leadership roles sooner than peers who don't.

But there's a catch. As individual scholars soar through the academic ranks, science as a whole shrinks its curiosity. AI-heavy research covers less novel ground, clusters around the same data-rich problems, and sparks less follow-on engagement between studies.

The findings highlight a tension between personal career advancement and collective scientific progress, as tools such as ChatGPT and

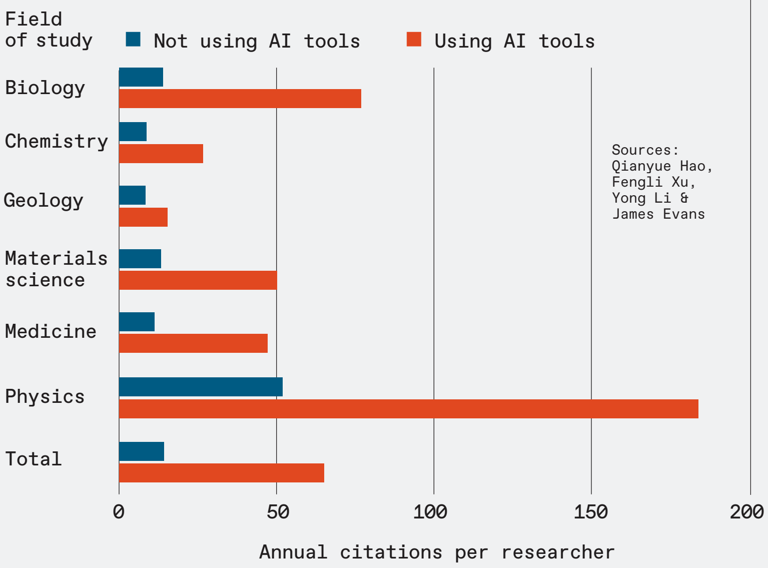

Sources: Qianyue Hao, Fengli Xu, Yong Li & James Evans

Across multiple fields, researchers who used AI tools in producing their research papers saw more annual citations than those who did not.

Researchers who used AI tools generally saw a decrease in "knowledge extent," which is functionally a decline in how much novel research space was covered in their work.

Sources: Qianyue Hao, Fengli Xu, Yong Li & James Evans

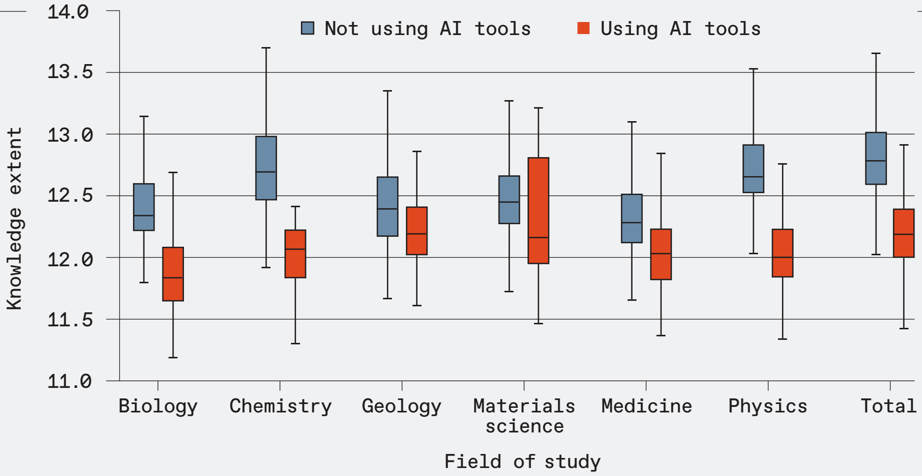

AlphaFold seem to reward speed and scale—but not surprise.

"You have this conflict between individual incentives and science as a whole," says James Evans, a sociologist at the University of Chicago who co-led the study.

And as more researchers pile onto the same scientific bandwagons, some experts worry about a feedback loop of conformity and declining originality. "This is very problematic," says Luís Nunes Amaral, a physicist who studies complex systems at Northwestern University, in Evanston, Ill., and who was not involved in the study. "We are digging the same hole deeper and deeper."

For Evans, the tension between efficiency and exploration is familiar terrain. He has spent more than a decade using massive publication and citation datasets to quantify how ideas spread, stall, and sometimes converge.

In 2008, he showed that the shift to online publishing and search made scientists more likely to read and cite the same highly visible papers, accelerating the dissemination of new ideas but narrowing the range of ideas in circulation. Later work detailed how career incentives quietly steer scientists toward safer, more crowded questions rather than riskier, original ones.

Another study from Evans tracked how large fields tend to slow their rate of conceptual innovation over time, even as the volume of papers explodes. And more recently, he has begun turning the same quantitative lens on AI itself, examining how algorithms reshape collective attention, discovery, and the organization of knowledge.

To quantify the AI effect, Evans and his collaborators from the Beijing National Research Center for Information Science and Technology trained a natural-language processing model to identify AI-augmented research across six natural science disciplines. They published their findings on 14 January in the journal *Nature*.

> "We are digging the same hole deeper and deeper."
>
> -LUÍS NUNES AMARAL, NORTHWESTERN UNIVERSITY

Their dataset included 41.3 million English-language papers published between 1980 and 2025 in biology, chemistry, physics, medicine, materials science, and geology. (They excluded fields such as computer science and mathematics that focus on developing AI methods themselves.) Evans's group compared roughly 311,000 papers that incorporated AI in some way—through the use of neural networks or large language models, for example—with millions of others that did not.

The researchers then traced the careers of individual scientists, examined how their papers accumulated attention, and zoomed out to consider how entire fields clustered or dispersed intellectually over time.

The results revealed a striking tradeoff. Scientists who adopt AI on average publish three times as many papers, receive nearly five times as many citations, and become team leaders a year or two earlier than those who do not.

But when those papers are mapped in a high-dimensional "knowledge space," AI-heavy research occupies a smaller intellectual footprint, clusters more tightly around popular, data-rich problems, and generates weaker networks of follow-on engagement between studies.

The pattern held across decades of AI development, spanning early machine learning, the rise of deep learning, and the current wave of generative AI. "If anything," Evans says, "it's intensifying."

Intellectual narrowing isn't the only unintended consequence. With automated tools making it easier to mass-produce manuscripts and conference submissions, journal editors and meeting organizers have witnessed a surge in low-quality and fraudulent papers, often produced at industrial scale.

"We've become so obsessed with the number of papers [that scientists publish] that we are not thinking about what it is that we are researching—and in what ways that contributes to a better understanding of reality, of health, and of the natural world," says Northwestern's Nunes Amaral.

At stake is the trajectory of scientific inquiry. "Certain types of questions are more amenable to AI tools," says Catherine Shea, who researches organizational behavior at Carnegie Mellon University's Tepper School of Business, in Pittsburgh, and who also was not involved in the research. In an academic environment in which papers are the main currency of success, researchers naturally gravitate toward the problems that are easiest for these tools to crank through and turn into publishable results. "It just becomes this self-reinforcing loop over time," she adds.

Models trained on abundant existing data excel at optimizing well-defined problems: predicting protein structures, classifying images, extracting patterns from massive datasets. Some systems have also begun to propose new hypotheses and research directions—a glimpse of what some call an "AI co-scientist."

But such systems—and the scientists who rely on them—are unlikely to venture into poorly mapped territories where data are scarce and questions are messier, Evans says, unless they are deliberately designed and given an incentive to do so. The danger is not that science slows down, but that it becomes more homogeneous—a high-speed version of the same narrowing Evans first documented when search engines replaced library stacks.

"This is a really scary paper to think about in terms of how the second- and third-order effects of using AI in science play out," says Shea.

Whether this trend persists may depend on how the next generation of AI tools is built and deployed across scientific workflows.

In a paper published in December, Bowen Zhou and his colleagues at the Shanghai Artificial Intelligence Laboratory, in China, argued that the application of AI in science remains fragmented. Data, computation, and hypothesis-generation tools are often deployed in a siloed and task-specific fashion. This limits knowledge transfer and blunts transformative discovery. But when those elements are integrated, AI-for-science systems help expand scientific discovery, says Zhou, a machine learning researcher who previously served as chief scientist of the IBM Watson group.

Evans doesn't think the problem is baked into the algorithmic design of AI. What matters most, he says, is overhauling the reward structures that shape what scientists work on in the first place.

For now, Evans says, the challenge is to deliberately redirect how AI is used and rewarded in science: "In some sense, we haven't fundamentally invested in the real value proposition of AI for science, which is asking what it might allow us to do that we haven't done before." ■

ENERGY

Airplane Beams Power to Ground Receiver During Flight › The demo tees up a new space-based solar power scheme

BY ANDREW MOSEMAN

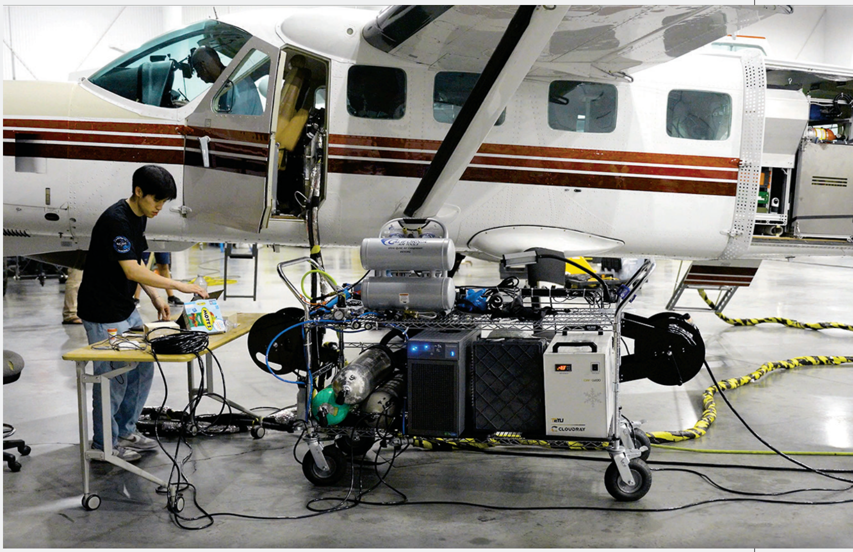

Michael Liang, an engineer at Overview Energy, works on the start-up's power-beaming equipment.

On a blustery November day, a Cessna turboprop flew over Pennsylvania at 5,000 meters, in crosswinds of up to 70 knots—about 130 kilometers per hour, nearly as fast as the little plane was flying. But the bumpy conditions didn't thwart its mission: to wirelessly beam power down to receivers on the ground as it flew over them.

The test flight marked the first time power has been beamed from a moving aircraft. It was conducted by the Ashburn, Va.–based startup Overview Energy, which emerged from stealth mode in December by announcing the feat.

ARTIFICIAL INTELLIGENCE

Machine Learning System Feels Your Pain

WHILE UNDERGOING SURGERIES with local anesthesia, some patients (such as infants or people with dementia) may have difficulty communicating whether they're feeling any pain. Researchers in Germany have now developed a contactless method using cameras and machine learning to estimate pain levels.

A machine learning model extracts data on facial expressions from a camera monitoring a patient. Heart-rate data can also be collected via remote photoplethysmography, in which light reflecting off a person's body is used to detect changes in blood volume.

Most existing pain-prediction algorithms are trained using very short, unrealistically "clean" video clips. But in actual surgical procedures, a patient's face may be obscured by a physician, poor lighting, or a piece of equipment. The new model was trained using longer, unedited video clips that reflect more realistic clinical situations.

As a result, the model's accuracy is relatively lower, at 45 percent. However, it theoretically could be more applicable in real-world settings than existing algorithms, and Reichard notes that using more sophisticated machine learning methods could further improve accuracy.

The study was published in *IEEE Open Journal of Engineering in Medicine and Biology*.

— By Michelle Hampson

But the greater purpose of the flight was to demonstrate the feasibility of a much grander ambition—to beam power from space to Earth. Overview plans to launch satellites into geosynchronous orbit (GEO) to collect unfiltered solar energy somewhere the sun never sets, and then beam this abundance back to humanity. The solar energy would be transferred as near-infrared waves and received by existing solar panels on the ground.

The far-flung strategy, known as space-based solar power, has become the subject of both daydreaming and serious research over the past decade. Caltech's Space Solar Power Project launched a demonstration mission in 2023 that transferred power, via microwaves, from transmitters to receivers aboard a satellite. And terrestrial power beaming is coming along too. The U.S. Defense Advanced Research Projects Agency (DARPA) in July 2025 set a new record for wirelessly transmitting power: 800 watts over 8.6 km for 30 seconds using a laser beam.

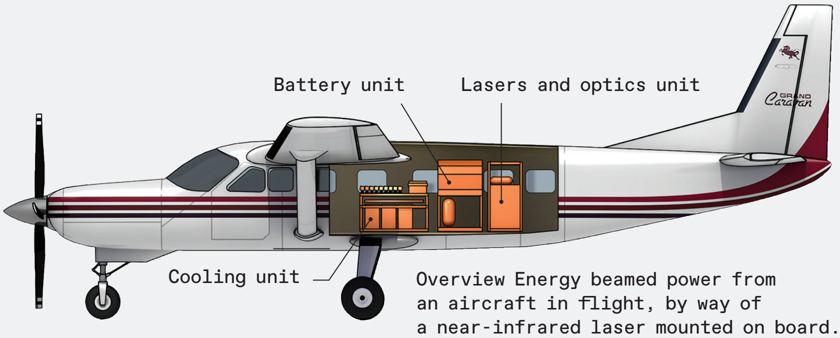

Overview Energy beamed power from an aircraft in flight, by way of a near-infrared laser mounted on board.

But until November, no one had actively beamed power from a moving platform to a ground receiver.

Overview's test transferred only a sprinkling of power, but it did so with the same components and techniques that the company plans to send to space. "Not only is it the first optical power beaming from a moving platform at any substantial range or power," says Overview CEO Marc Berte, "but also it's the first time anyone's really done a power-beaming thing where it's all of the functional pieces all working together. It's the same methodology and function that we will take to space and scale up in the long term."

The approach was compelling enough that power-beaming expert Paul Jaffe left his job as a program manager at DARPA to join the company as head of systems engineering. Prior to DARPA, Jaffe spent three decades with the U.S. Naval Research Laboratory.

It was hearing Berte explain Overview's plan at a conference that helped to convince Jaffe to take a chance on the startup. "This actually sounds like it could work," Jaffe remembers thinking at the time. "It really seems like it gets around a lot of the showstoppers for a lot of the other concepts. I remember coming home and telling my wife that I almost felt like the problem had been solved. So I thought: Should [I] do something which is almost unheard of—to leave in the middle of being a DARPA program manager—to try to do something else?"

For Jaffe, the most compelling reason was Overview's solution for space-based solar's power-density problem. A beam with lower power density is safer because it's not blasting too much concentrated energy onto a single spot on Earth's surface, but it's less efficient for the task of delivering usable solar energy. A higher-density beam does the job better, but then the researchers must engineer some way to maintain safety.

Many researchers have settled on microwaves as their beam of choice for wireless power. But, in addition to the safety concerns about shooting such intense radiation at Earth, Jaffe says there's another problem: Microwaves include part of what he calls the "beachfront property" of the electromagnetic spectrum—a range from 2 to 20 gigahertz that is set aside for many other applications, such as 5G cellular networks.

"The fact is," Jaffe says, "if you somehow magically had a fully operational solar power satellite that used microwave power transmission in orbit

LEFT: GETTY IMAGES; RIGHT: HUANYU ZHOU, HYUN-WOOK KIM, ET AL.

Illustration by John MacNeill

today—and a multikilometer-scale microwave power satellite receiver on the ground magically in place today—you could not turn it on because the spectrum is not allocated to do this kind of transmission."

Instead, Overview plans to use wide-beam near-infrared waves. Existing utility-scale solar farms could receive the beamed energy just like they receive the sun's energy during daylight hours. So "your receivers are already built," Berte says. The next major step is a prototype demonstrator for low Earth orbit, after which he hopes to have GEO satellites beaming megawatts of power by 2030 and gigawatts by later that decade.

Plenty of doubts about the feasibility of space-based power abound. It is an exotic technology with much left to prove, including the ability to survive damage by orbital debris and the exorbitant cost of launching the power stations. (Overview's satellite will be built on Earth in a folded configuration, and it will unfold after it's brought to orbit, according to the company.)

"Getting down the cost per unit mass for launch is a big deal," Jaffe says. "Then, it just becomes a question of increasing the specific power. A lot of the technologies we're working on at Overview are squarely focused on that." ■

CONSUMER ELECTRONICS

Stretchable OLEDs Just Got a Huge Upgrade › 2D material lets displays stretch 200 percent and stay bright

BY PERRI THALER

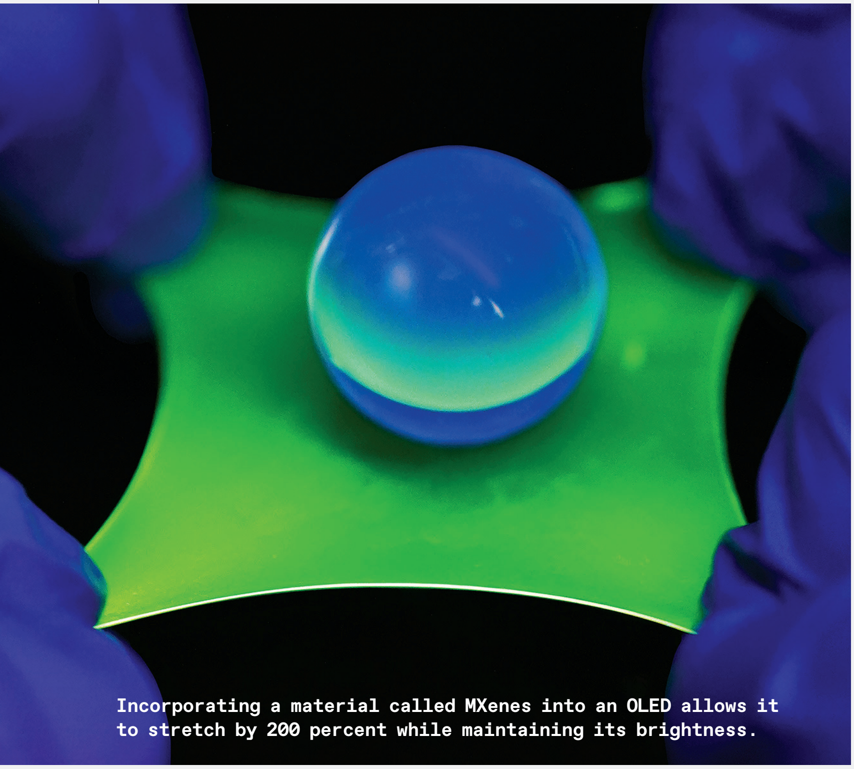

Incorporating a material called MXenes into an OLED allows it to stretch by 200 percent while maintaining its brightness.

Engineers developing wearable displays have long sought OLEDs that can bend, twist, and stretch while maintaining a bright and stable light. These displays could be integrated into new classes of devices—woven into clothing fabric, for example, to show real-time information like a runner's speed or heart rate, without breaking or dimming.

But engineers have always encountered a trade-off: The more you stretch these materials, the dimmer they become. A group co-led by Yury Gogotsi, a materials scientist at Drexel University, in Philadelphia, has found a way around the problem by employing a special class of materials called MXenes—which Gogotsi helped discover in 2011. With MXenes (pronounced "max-enes"), the team developed an OLED that can stretch to twice its original size while keeping a steady glow. The OLED also converts electricity into light more efficiently than any stretchable OLED before it.

Gogotsi didn't have much experience with OLEDs when, about five years ago, he teamed up with Tae-Woo Lee, a materials scientist at Seoul National University. Lee was driven to develop better flexible OLEDs by the ever-increasing use of flexible electronics like foldable phones.

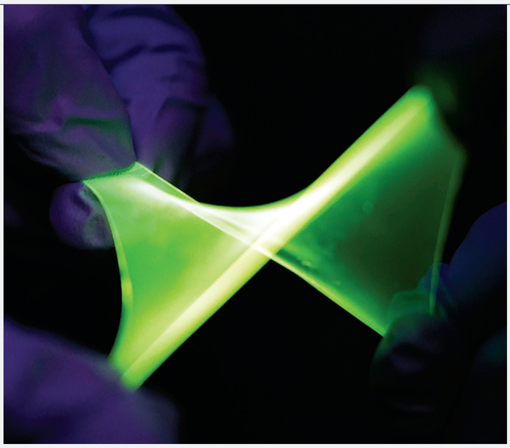

The new OLED material can also be twisted and bent without breaking.

Traditionally, displays are built from multiple stacked layers sandwiched between glass. At the base, a cathode supplies electrons that enter into adjacent organic layers, which are designed to conduct this charge efficiently. As the electrons move through these layers, they meet a positive charge produced by an indium tin oxide (ITO) film. The moment these charges combine, the organic material releases energy as light, creating the illuminated pixels that make up an image. The entire structure is sealed with a glass layer on top.

The ITO film—adhered to the glass—serves as a transparent anode, allowing current to pass into the organic layers without blocking the generated light. "But it's brittle. It's ceramic, basically," Gogotsi says, so while it works well for flat surfaces, it can't be bent.

In an effort to replace the ITO film with a stretchable alternative, Gogotsi's students started by creating a transparent, conducting film out of a MXene, a type of ultrathin and flexible material with metal-like conductivity. The material is unique in its inherent ability to bend because it's made from many two-dimensional sheets that can slide relative to one another without breaking. The film—only 10 nanometers thick—"appeared to be this perfect replacement for ITO," Gogotsi says.

Through experimentation, Gogotsi and Lee's shared team found that a mix of the MXene and silver nanowire would actually stretch the most while maintaining stability. "We were able to double the size, achieving 200 percent stretching without losing performance," Gogotsi says.

And by creating a more-efficient contact between the topmost light-emitting organic layer and the new MXene film, the team increased the film's flexibility and brightness by almost an order of magnitude compared with that of ITO.

That's because, unlike ITO, the surface of MXenes can be chemically adjusted to make it easier for electrons to move from the electrode into the light-emitting layer. This more efficient electron flow significantly increases the brightness of the display, as evidenced by a 17 percent external quantum efficiency—a measure of how

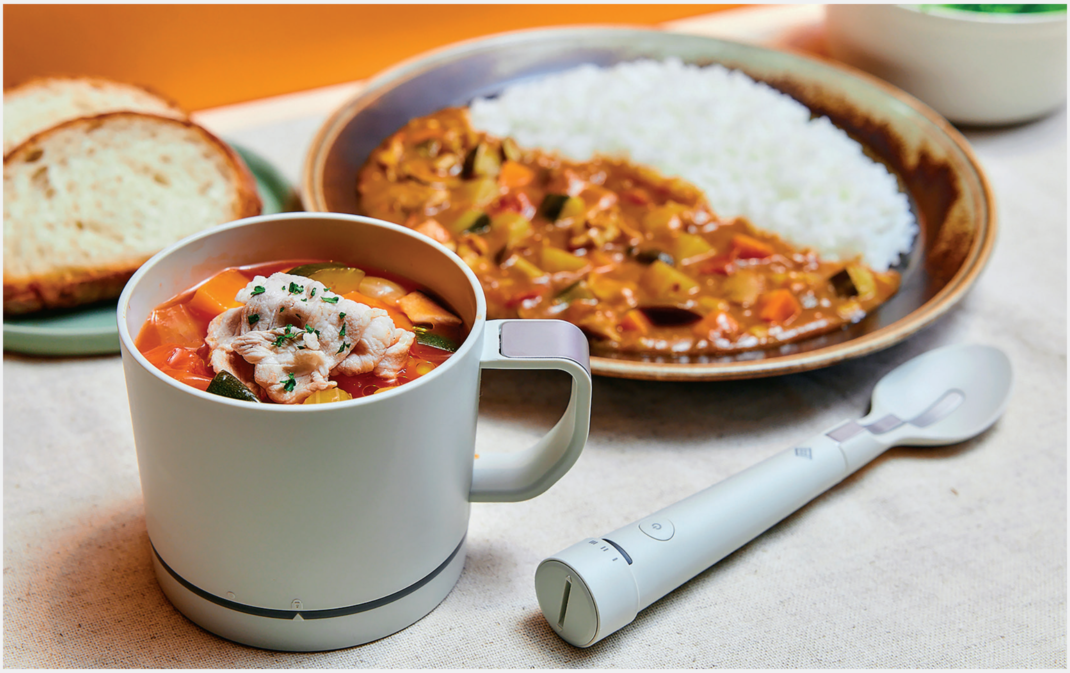

THE JAPANESE BEVERAGE COMPANY KIRIN developed the Electric Salt Spoon and Cup to make foods taste saltier without adding more sodium. The spoon and cup apply a weak electric current to the food they hold to drive the existing sodium ions toward the eater's tongue. Another, similar attempt by Kirin and Meiji University researcher Homei Miyashita, this time using a chopstick, required a wristband component to provide the electrical stimulation. Kirin's spoon and cup, which the company is launching in Asian markets outside Japan this year, also come with three levels of perceived saltiness. **—Novid Parsi**

TOP: HUANYU ZHOU, HYUN-WOOK KIM, ET AL.; BOTTOM: KIRIN

efficiently a device turns electricity into light—which the group claims is a record for stretchable OLEDs.

"Achieving those numbers in intrinsically stretchable OLEDs under substantial stretching is quite significant," says Seunghyup Yoo, who runs the Integrated Organic Electronics Lab at KAIST, formerly known as the Korea Advanced Institute of Science and Technology, in Daejeon, South Korea. An external quantum efficiency of 20 percent is an important benchmark for this kind of device because it is the upper limit of efficiency dictated by the physics, Yoo explains.

To increase illumination, the researchers went beyond working with MXene. Lee's group developed two additional organic layers to add into the middle of their OLED—one that directs positive charges to the light-emitting layer, ensuring that electricity is used more efficiently, and one that recycles wasted energy that would normally be lost, boosting overall brightness.

Together, the MXene layer and two organic layers allow for a notably bright and stable OLED, even when stretched. Gogotsi thinks the resulting OLED is "very successful" because it combines both brightness and stretchability, while historically, engineers have been able to achieve only one or the other.

"The performance that they are able to achieve in this work is an important advancement," says Sihong Wang, a molecular engineer at the University of Chicago who also develops stretchable OLED materials. Wang also notes that the 200 percent stretchability that Gogotsi's group attained is beyond what's needed for wearable applications.

A stretchable OLED that maintains its brightness has uses in many settings besides wearable devices, including communications, industrial manufacturing, and robotics, Gogotsi says, although he's most excited about its adoption in health-monitoring devices. He sees a near future in which displays for diagnostics and treatment become embedded in clothing or as "epidermal electronics."

Before these displays can come to market, however, stability issues inherent to all stretchable OLEDs need to be solved, Wang says. Today's materials are not able to sustain light emissions long enough to maintain images.

Finding housings to protect the OLEDs from damage is also a problem. "You need a stretchable encapsulation material that can protect the central device without allowing oxygen and moisture to permeate," Wang says.

Yoo agrees: He says it's a tough problem to solve because the best protective layers are rigid and not very stretchable. He notes that developing stretchable displays that don't distort images are another challenge in the path to commercialization.

Regardless, Gogotsi is excited about the future of stretchable OLEDs. "We started with computers occupying the room, then moved to our desktops, then to laptops, then we got smartphones and iPads, but still we carry stuff with us," he says. "Flexible displays can be on the sleeve of your jacket. They can be rolled into a tube or folded and put in your pocket. They can be everywhere." ■

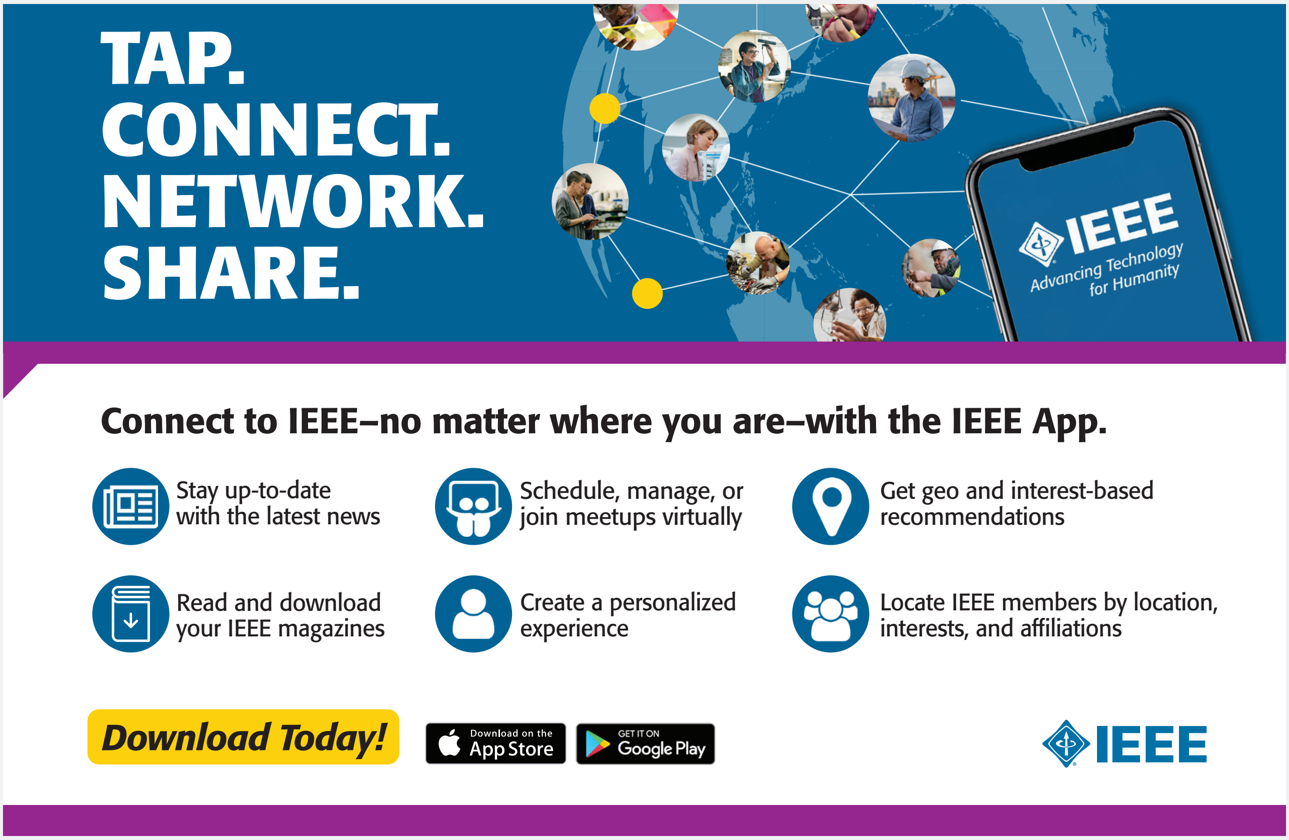
TAP.
CONNECT.
NETWORK.
SHARE.

Connect to IEEE–no matter where you are–with the IEEE App.

Stay up-to-date with the latest news

Schedule, manage, or join meetups virtually

Get geo and interest-based recommendations

Read and download your IEEE magazines

Create a personalized experience

Locate IEEE members by location, interests, and affiliations

Download Today!

Download on the App Store

GET IT ON Google Play

IEEE

The Data

A Lofty Idea for AI Data Centers

What's the difference between a stupid idea and a brilliant one? Sometimes, it just comes down to resources. Practically unlimited funds, like limitless thrust, can get even a mad idea off the ground.

And so it might be for the concept of putting AI data centers in orbit. In a rare moment of unalloyed agreement, some of the richest and most powerful men in technology are staunchly backing the idea. The group includes Elon Musk, Jeff Bezos, Jensen Huang, Sam Altman, and Google CEO Sundar Pichai. In all likelihood, hundreds of people are now working on the concept of space data centers at the firms directly or indirectly controlled by these men—SpaceX, Starlink, Tesla, Amazon, Blue Origin, Nvidia, OpenAI, and Google, among others.

So how much would it cost to start training large language models in space? Probably the best accounting is one created by aerospace engineer Andrew McCalip. His exhaustive, detailed analysis includes interactive sliders that let you compare costs for space-based and terrestrial data centers in the range of 1 to 100 gigawatts. One-gigawatt data centers are being built now on terra firma, and Meta has announced plans for a 5-GW facility, with anticipated completion some time after 2030.

In an interview, McCalip says his initial rough calculations a few years ago suggested that data centers in space would cost in the range of 7 to 10 times more, per gigawatt of capacity, than their terrestrial counterparts. "It just wasn't practical," he says. "Not even close." But when Elon Musk began publicly backing the idea, McCalip revisited the numbers using publicly available information about Starlink's and Tesla's technologies and capabilities.

That changed the picture substantially. The figures in his online analysis assume an orbital network of data-center satellites that borrows heavily from Musk's tech treasure chest—"essentially…you just start putting some radiation-resistant ASIC chips on the Starlink fleet and you start growing edge capacity organically on the Starlink fleet," McCalip says. The network would rely on the kind of watt-efficient GPU architecture used in Teslas for self-driving, he adds. "You start dropping those onto the backs of Starlinks. You can slowly grow this out, and this would be approximately the performance that you would get."

Bottom line, with some solid but not necessarily heroic engineering, the cost of an orbital data center could be as low as three times that of the comparable terrestrial one. That differential, while still high, at least nudges the concept out of the instantly dismissible category. "I have my particular views, but I want the data to speak for itself," McCalip says.

For this illustration, we picked a configuration with an aggregate 1 GW of capacity. The network would consist of some 4,400 satellites, each of which would be outfitted with a 1,024-square-meter solar array that generates 240 kilowatts. The data center on that satellite, powered by the array, might have at least 175 GPUs; McCalip notes that a popular GPU rack, Nvidia's NVL72 line, has 72 GPUs and requires 120 to 140 kW.

The total cost of the satellite network would be around US $51 billion, including launch and five years of operational expenses; a comparable terrestrial system would cost about $16 billion over the same period.

Stupid? Not stupid? You decide. ■

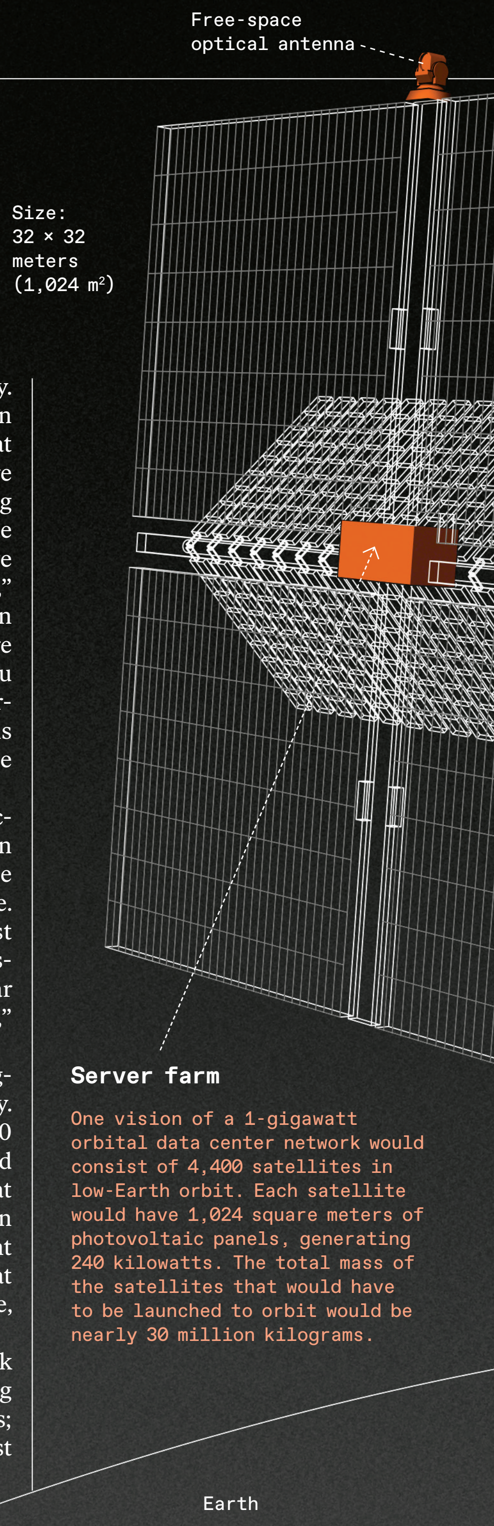
Free-space optical antenna
Size: 32 × 32 meters (1,024 m^2)
Earth

Server farm

One vision of a 1-gigawatt orbital data center network would consist of 4,400 satellites in low-Earth orbit. Each satellite would have 1,024 square meters of photovoltaic panels, generating 240 kilowatts. The total mass of the satellites that would have to be launched to orbit would be nearly 30 million kilograms.

Illustration by John MacNeill

BY GLENN ZORPETTE

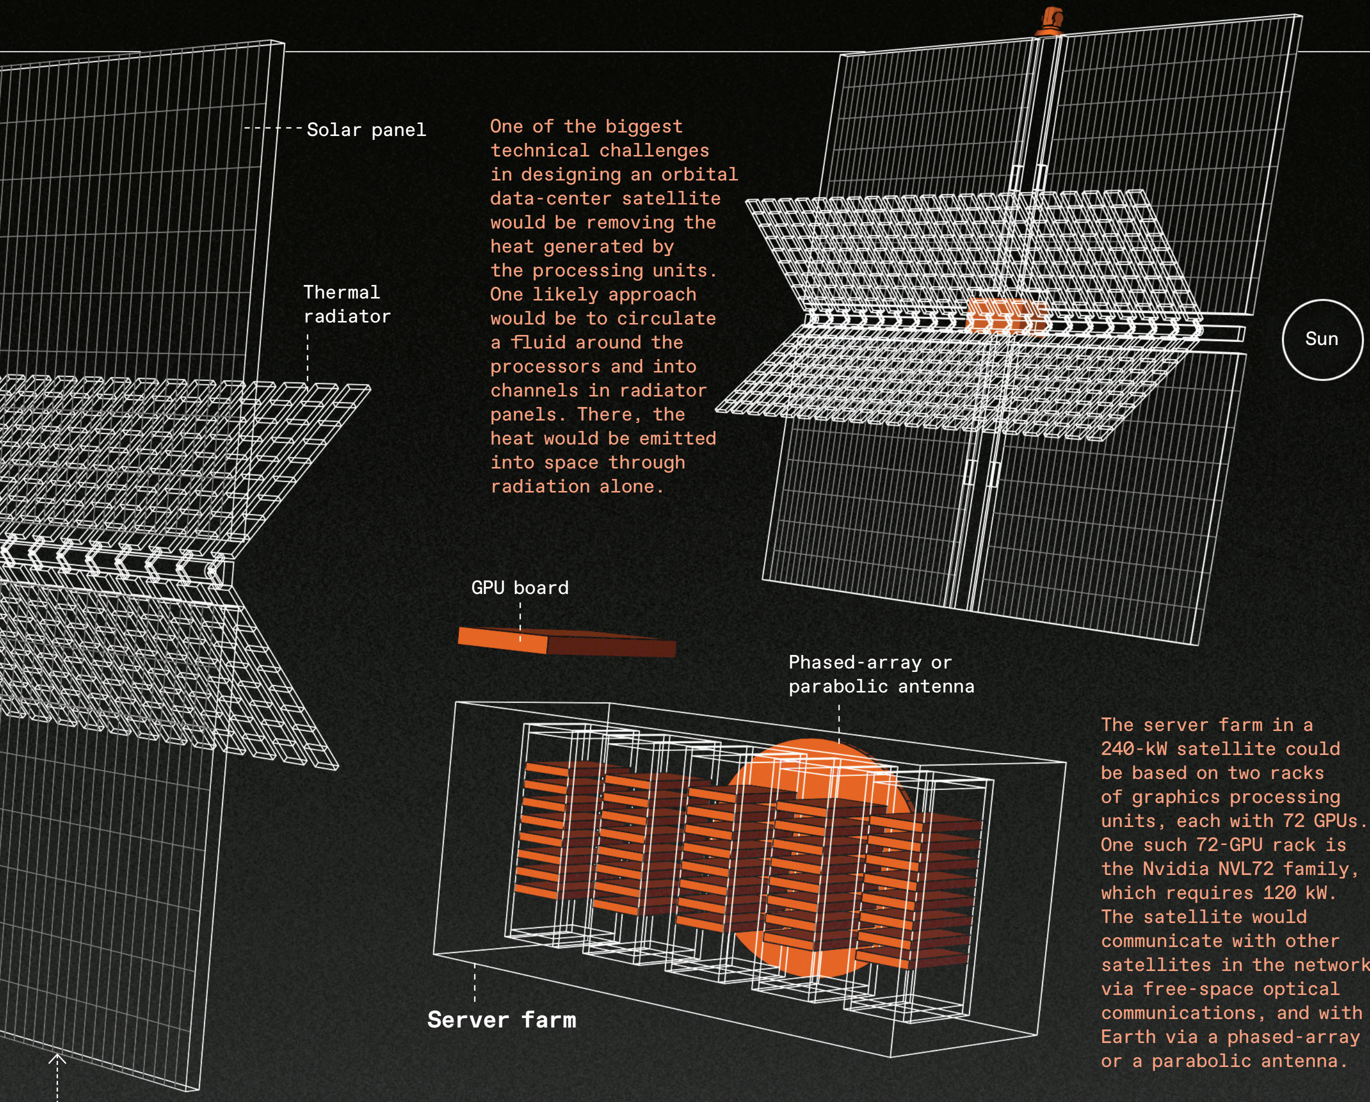

One of the biggest technical challenges in designing an orbital data-center satellite would be removing the heat generated by the processing units. One likely approach would be to circulate a fluid around the processors and into channels in radiator panels. There, the heat would be emitted into space through radiation alone.

The server farm in a 240-kW satellite could be based on two racks of graphics processing units, each with 72 GPUs. One such 72-GPU rack is the Nvidia NVL72 family, which requires 120 kW. The satellite would communicate with other satellites in the network via free-space optical communications, and with Earth via a phased-array or a parabolic antenna.

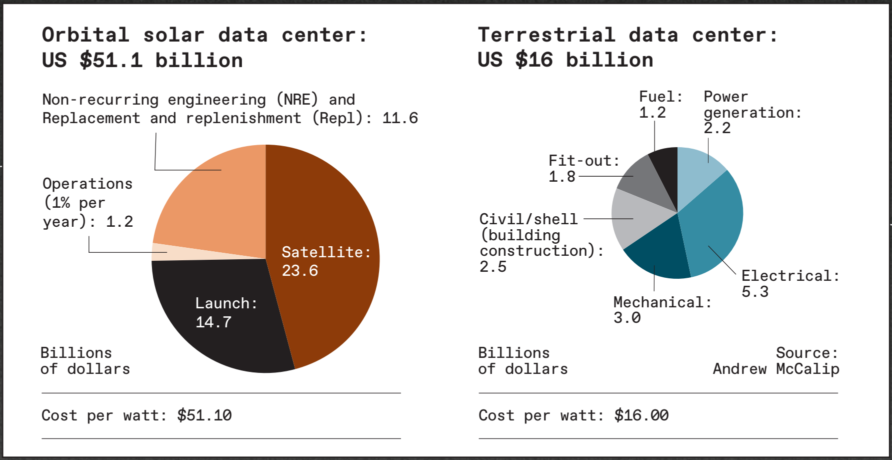

Source: Andrew McCalip

Likely costs to design, build, and launch a 1-GW orbital data center, based on a network of some 4,400 satellites and including operating costs over a five-year period, would exceed US $50 billion. That's about three times the cost of a 1-GW data center on Earth, including five years of operation.

TECH TO TINKER WITH

Hands On

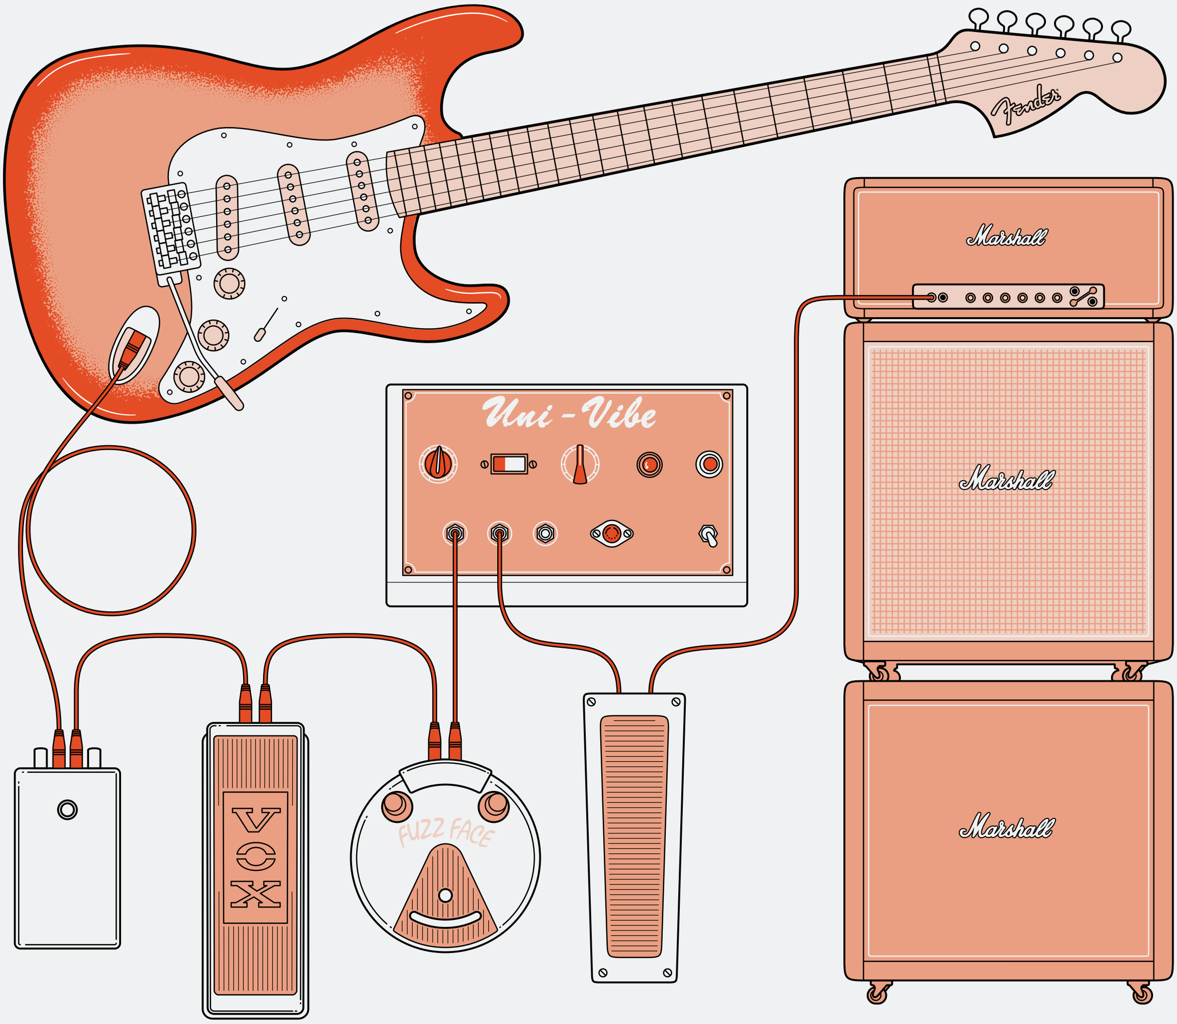

Jimi Hendrix used a chain of components to modulate sound, controlling a feedback loop by positioning the guitar with respect to his amplifier's speaker.

Jimi Hendrix, Systems Engineer › His iconic sound came from precise modulation and feedback

BY ROHAN S. PURANIK

3 February 1967 is a day that belongs in the annals of music history. It's the day that Jimi Hendrix entered London's Olympic Studios to record a song using a new component. The song was "Purple Haze," and the component was the Octavia guitar pedal, created for Hendrix by sound engineer Roger Mayer. The pedal was a key element of a complex chain of analog elements responsible for the final sound, including the acoustics of

Illustrations by James Provost

the studio room itself. When they sent the tapes for remastering in the United States, the sounds on it were so novel that they included an accompanying note explaining that the distortion at the end was not malfunction but intention. A few months later, Hendrix would deliver his legendary electric guitar performance at the Monterey International Pop Festival.

"Purple Haze" firmly established that an electric guitar can be used not just as a stringed instrument with built-in pickups for convenient sound amplification, but also as a full-blown wave synthesizer whose output can be manipulated at will. Modern guitarists can reproduce Hendrix's chain using separate plug-ins in digital audio workstation software, but the magic often disappears when everything is buffered and quantized. I wanted to find out if a more systematic approach could do a better job and provide insights into how Hendrix created his groundbreaking sound.

My fascination with Hendrix's Olympic Studios performance arose because there is a "Hendrix was an alien" narrative surrounding his musical innovation—that his music appeared more or less out of nowhere. I wanted to replace that narrative with an engineering-driven account that's inspectable and reproducible—plots, models, and a signal chain from the guitar through the pedals that you can probe stage by stage.

Although I work mostly in the digital domain as an edge-computing architect in my day job, I knew that analog circuit simulations would be the key to going deeper.

My first step was to look at the challenges Hendrix was trying to address. Before the 1930s, guitars were too quiet for large ensembles. Electromagnetic pickups—coils of wire wrapped around magnets that detect the vibrations of metal strings—fixed the loudness problem. But they left a new one: the *envelope*, which specifies how the amplitude of a note varies as it's played on an instrument, starting with a rising initial *attack*, followed by a falling *decay*, and then

Each effects pedal in Hendrix's chain contributed to enhancing the electric guitar beyond its intrinsic limits. A selection of plots from the full-circuit ngspice analysis shows how the Fuzz Face turns a sinusoid signal from a string into an almost square wave [top]; how the Octavia pedal inverts half the input waveform to double its frequency [second from top]; how the wah-wah pedal acts as band-pass filter [second from bottom]; and how the Uni-Vibe pedal introduces selective phase shifts to color the sound [bottom].

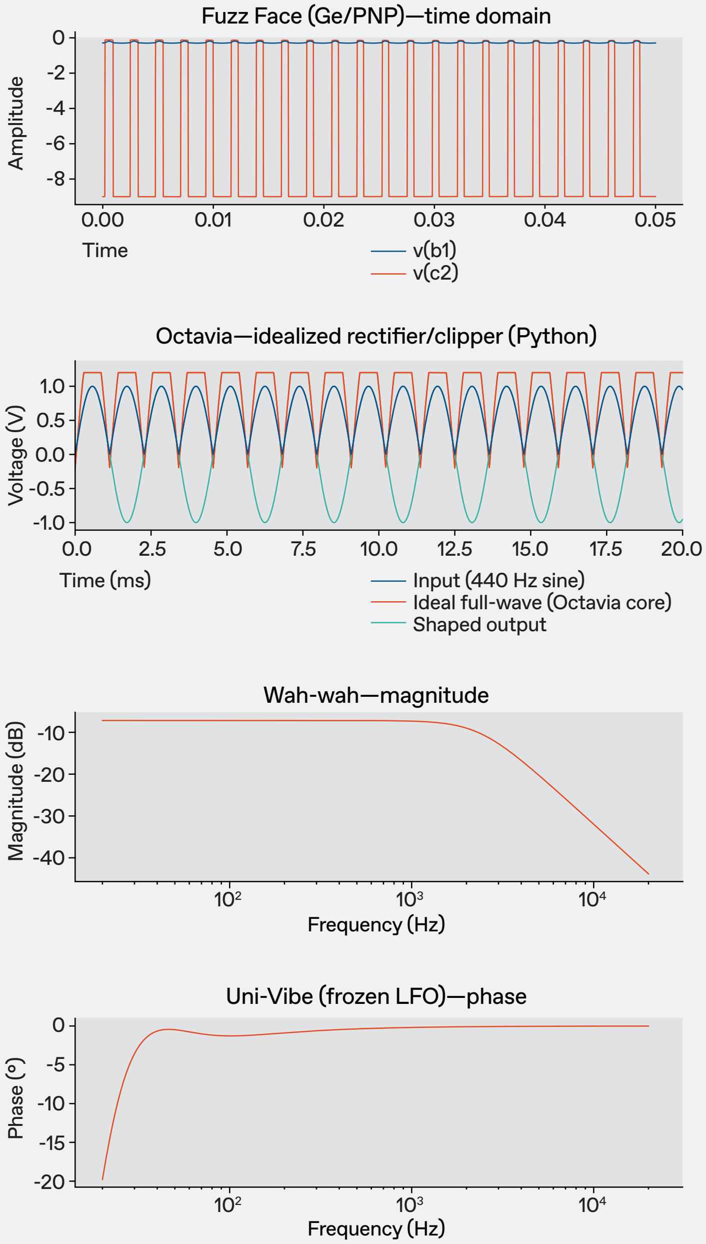

The Jimi Hendrix Experience [(left to right) Mitch Mitchell, Jimi Hendrix, Noel Redding] pose during construction of a recording studio in New York City.

FRED W. MCDARRAH/GETTY IMAGES

any *sustain* of the note after that. Electric guitars attack hard, decay fast, and don't sustain like bowed strings or organs. Early manufacturers tried to modify the electric guitar's characteristics by using hollow bodies fitted with magnetic pickups, but the instrument still barked more than it sang.

Hendrix's mission was to reshape both the electric guitar's envelope and its tone until it could feel like a human voice. He tackled the guitar's constraints by augmenting it. His solution was essentially a modular analog signal chain driven not by knobs but by hands, feet, gain staging, and physical movement in a feedback field.

Hendrix's setups are well documented: Set lists, studio logs, and interviews with Mayer and Eddie Kramer, then the lead engineer at Olympic Studios, fill in the details. The signal chain for "Purple Haze" consisted of a set of pedals—a Fuzz Face, the Octavia, and a wah-wah—plus a Marshall 100-watt amplifier stack, with the guitar and room acoustics closing a feedback loop that Hendrix tuned with his own body. Later, Hendrix would also incorporate a Uni-Vibe pedal for many of his tracks. All the pedals were commercial models except for the Octavia, which Mayer built to produce a distorted signal an octave higher than its input.

I obtained the schematics for each of these elements and their accepted parameter ranges, and converted them into netlists that the ngspice can process (ngspice is an open-source implementation of the circuit simulator Spice). The Fuzz Face pedal came in two variants, using germanium or silicon transistors, so I created models for both. In my models, Hendrix's guitar pickups had a resistance of 6 kiloohms and an inductance of 2.5 henrys with a realistic cable capacitance.

I chained the ngpsice circuit simulations together into a whole using a script, and I produced data-plot and sample sound outputs with Python scripts. All of the ngspice and scripts are available in my GitHub repository at github.com/nahorov/Hendrix-Systems-Lab, with instructions on how to reproduce my simulations.

Plotting the signal at different points in the chain with different parameters reveals how Hendrix configured and manipulated the nonlinear complexities of the system as a whole to reach his expressive goals.

A few highlights: First, the Fuzz Face is a two-transistor feedback amplifier that turns a gentle sinusoid signal into an almost binary "fuzzy" output. The interesting behavior emerges when the guitar's volume is reduced. Because the pedal's input impedance is very low (about 20 kΩ), the pickups interact directly with the pedal circuit. Reducing amplitude restores a sinusoidal shape—producing the famous "cleanup effect" that was a hallmark of Hendrix's sound, where the fuzz drops in and out as desired while he played.

Second, the Octavia pedal used a rectifier, which normally converts alternating to direct current. Mayer realized that a rectifier effectively flips each trough of a waveform into a peak, doubling the number of peaks per second. The result is an apparent doubling of frequency—a bloom of second-harmonic content that the ear hears as a bright octave above the fundamental.

Third, the wah-wah pedal is a bandpass filter: Frequency plots show the center frequency sweeping from roughly 300 hertz to 2 kilohertz. Hendrix used it to make the guitar "talk" with vowel sounds, most iconically on "Voodoo Child (Slight Return)."

Fourth, the Uni-Vibe cascades four phase-shift sections controlled by photoresistors. In circuit terms, it's a low-frequency oscillator modulating a variable-phase network; in musical terms it's motion and air.

Finally, the whole chain became a closed loop by driving the Marshall amplifier near saturation, which among other things extends the sustain. In a reflective room, the guitar strings couple acoustically to the speakers—move a few centimeters and you shift from one stable feedback mode to another. To an engineer, this is a gain-controlled acoustic feedback system. To Hendrix, it was part of the instrument. He learned to tune oscillation with distance and angle, shaping sirens, bombs, and harmonics by walking the edge of instability.

Hendrix didn't speak in decibels and ohm values, but he collaborated with engineers who did—Mayer and Kramer—and iterated fast as a systems engineer. Reframing Hendrix as an engineer doesn't diminish the art. It explains how one person, in under four years as a bandleader, could pull the electric guitar toward its full potential by systematically augmenting the instrument's shortcomings for maximum expression. ■

BY GWENDOLYN RAK

Careers

Though Alan DeKok's work goes unnoticed by most, his FreeRADIUS software is used every day to authenticate users.

KAMARA MOROZUK

Alan DeKok

Serendipity led him from nuclear physics to a career in network security

When Alan DeKok began a side project in network security, he didn't expect to start a 27-year career. In fact, he didn't initially plan to work in computing at all.

DeKok studied nuclear physics before making the switch to a part of network computing that is foundational but—like nuclear physics—largely invisible to those not directly involved in the field. Eventually, a project he started as a hobby became a full-time job: maintaining one of the primary systems that helps keep the internet secure.

Today, he leads the FreeRADIUS Project, which he cofounded in the late 1990s to develop what is now the most widely used Remote Authentication Dial-In User Service (RADIUS) software. FreeRADIUS is an open-source server that provides back-end authentication for most major internet service providers. It's used by global financial institutions, Wi-Fi services like Eduroam, and Fortune 50 companies. DeKok is also CEO of InkBridge Networks, which maintains the server and provides support for the companies that use it.

Reflecting on nearly three decades of experience leading FreeRADIUS, DeKok says he became an expert in remote authentication "almost by accident," and the key to his career has largely been luck. "I really believe that it's preparing yourself for luck, being open to it, and having the skills to capitalize on it."

Employer:
InkBridge Networks

Occupation:
CEO

Education:
Bachelor's degree in physics, Carleton University; master's degree in physics, Carleton University

DeKok grew up on a farm outside of Ottawa growing strawberries and raspberries. "Sitting on a tractor in the heat is not particularly interesting," says DeKok, who was more interested in working with 8-bit computers than crops. As a student at Carleton University, in Ottawa, he found his way to physics because he was interested in math but preferred the practicality of science.

While pursuing a master's degree in physics, also at Carleton, he worked on a water-purification system for the Sudbury Neutrino Observatory, an underground observatory then being built at the bottom of a nickel mine. He would wake up at 4:30 in the morning to drive up to the site, descend 2 kilometers, then enter one of the world's deepest clean-room facilities to work on the project. The system managed to achieve one atom of impurity per cubic meter of water, "which is pretty insane," DeKok says.

But after his master's degree, DeKok decided to take a different route. Although he found nuclear physics interesting, he says he didn't see it as his life's work. Meanwhile, the Ph.D. students he knew were "fanatical about physics." He had kept up his computing skills through his education, which involved plenty of programming, and decided to look for jobs at computing companies. "I was out of physics. That was it."

Still, physics taught him valuable lessons. For one, "You have to understand the big picture," DeKok says. "The ability to tell the big-picture story in standards, for example, is extremely important." This skill helps DeKok explain to standards bodies how a protocol acts as one link in the entire chain of events that needs to occur when a user wants to access the internet.

He also learned that "methods are more important than knowledge." It's easy to look up information, but physics taught DeKok how to break down a problem into manageable pieces to come up with a solution. "When I was eventually working in the industry, the techniques that came naturally to me, coming out of physics, didn't seem to be taught as well to the people I knew in engineering," he says. "I could catch up very quickly."

In 1996, DeKok was hired as a software developer at a company called Gandalf, which made equipment for ISDN, a precursor to broadband that enabled digital transmission of data over telephone lines. Gandalf went under about a year later, and he joined CryptoCard, a company providing hardware devices for two-factor authentication.

While at CryptoCard, DeKok began spending more time working with a RADIUS server. When users want to connect to a network, RADIUS acts as a gatekeeper and verifies their identity and password, determines what they can access, and tracks sessions. DeKok moved on to a new company in 1999, but he didn't want to lose the networking skills he had developed. No open-source RADIUS servers were being actively developed at the time, and he saw a gap in the market.

"Almost by accident, I became one of the more senior people in the space"

Also in 1999, he started FreeRADIUS in his free time and it "gradually took over my life," DeKok says. He continued to work on the open-source software as a hobby for several years while bouncing around companies in California and France. "Almost by accident, I became one of the more senior people in the space. Then I doubled down on that and started the business." He founded NetworkRADIUS (now called InkBridge Networks) in 2008.

By that point, FreeRADIUS was already being used by 100 million people daily. The company now employs experts in Canada, France, and the United Kingdom who work together to support FreeRADIUS. "I'd say at least half of the people in the world get on the internet by being authenticated through my software," DeKok estimates. He attributes that growth largely to the software being open source. Initially a way to enter the market with little funding, going open source has allowed FreeRADIUS to compete with bigger companies as an industry-leading product.

Although the software is critical for maintaining secure networks, most people aren't aware of it because it works behind the scenes. DeKok is often met with surprise that it's still in use. He compares RADIUS to a building foundation: "You need it, but you never think about it until there's a crack in it."

Over the years, DeKok has maintained FreeRADIUS by continually making small fixes. Like using a ratcheting tool to make a change inch by inch, "you shouldn't underestimate that ratchet effect of tiny little fixes that add up over time," he says.

He's seen the project through minor patches and more significant fixes, like when researchers exposed a widespread vulnerability DeKok had been trying to fix since 1998. He also watched a would-be successor to the network protocol, Diameter, rise and fall in popularity in the 2000s and 2010s. (Diameter gained traction in mobile applications but has gradually been phased out in the shift to 5G.) Though Diameter offers improvements, RADIUS is far simpler and already widely implemented, giving it an edge, DeKok explains.

And he remains confident about its future. "People ask me, 'What's next for RADIUS?' I don't see it dying." Estimating that billions of dollars of equipment run RADIUS, he says, "It's never going to go away."

About his own career, DeKok says he plans to keep working on FreeRADIUS, exploring new markets and products. "I never expected to have a company and a lot of people working for me, my name on all kinds of standards, and customers all over the world. But it worked out that way." ■

BY DINA GENKINA

5 Questions

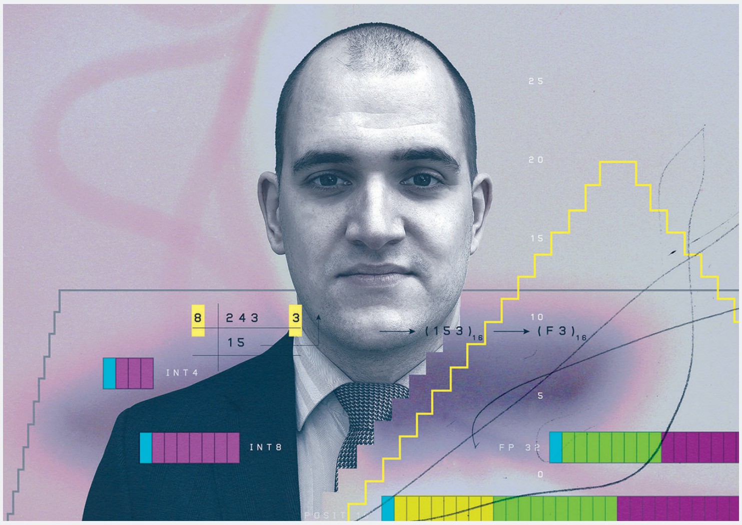

Laslo Hunhold

AI's new number formats don't work elsewhere

Artificial intelligence has driven an explosion of new number formats—the ways in which numbers are represented digitally. Engineers are looking at every possible way to save computation time and energy, including shortening the number of bits used to represent data. But what works for AI doesn't necessarily work for scientific computing, be it for computational physics, biology, fluid dynamics, or engineering simulations. *IEEE Spectrum* spoke with Laslo Hunhold, who recently joined Barcelona-based Openchip as an AI engineer, about developing a bespoke number format for scientific computing.

What makes number formats interesting to you?
Laslo Hunhold: I don't know another example of a field that so few are interested in but has such a high impact. If you make a number format that's 10 percent more [energy] efficient, it can translate to all applications being 10 percent more efficient, and you can save a lot of energy.

Why are there so many new number formats?
Hunhold: For decades, computer users had it really easy. They could just buy new systems every few years, and they would have performance benefits for free. But this hasn't been the case for the last 10 to 20 years. In computers, you have a certain number of bits used to represent a single number, and for years the default was 64 bits. And for AI, companies noticed that they don't need 64 bits for each number. So they had a strong incentive to go down to 16, 8, or even 2 bits [to save energy]. The problem is, the dominating standard for representing numbers in 64 bits is not well designed for lower bit counts. So in the AI field, they came up with new formats that are more tailored toward AI.

Why does AI need different number formats than scientific computing?
Hunhold: Scientific computing needs high dynamic range: You need very large numbers, or very small numbers, and high accuracy in both cases. The 64-bit standard has an excessive dynamic range, and it is many more bits than you need most of the time. It's different with AI. The numbers usually follow a specific distribution, and you don't need as much accuracy.

What makes a number format "good"?
Hunhold: You have infinite numbers but only finite bit representations. So you need to decide how you assign numbers. The most important part is to represent numbers that you're actually going to use. Because if you represent a number that you don't use, you've wasted a representation. The simplest thing to look at is the dynamic range. The next is distribution: How do you assign your bits to certain values? Do you have a uniform distribution, or something else? There are infinite possibilities.

What motivated you to introduce the takum number format?
Hunhold: Takums are based on posits. In posits, the numbers that get used more frequently can be represented with more density. But posits don't work for scientific computing, and this is a huge issue. They have a high density for [numbers close to one], which is great for AI, but the density falls off sharply once you look at larger or smaller values. People have been proposing dozens of number formats in the last few years, but takums are the only number format that's actually tailored for scientific computing. I found the dynamic range of values you use in scientific computations, if you look at all the fields, and designed takums such that when you take away bits, you don't reduce that dynamic range. ■

Laslo Hunhold is a senior AI accelerator engineer at Barcelona-based startup Openchip. He recently completed a Ph.D. in computer science from the University of Cologne, in Germany.

Photo-illustration by Stuart Bradford

AI HUNTS FOR THE NEXT

There's a crisis in particle physics. Researchers hope AI can help

BY MATTHEW HUTSON

This historic cloud-chamber image shows the spiral tracks of charged particles—an early, visual way physicists studied the subatomic world.
OMIKRON/SCIENCE SOURCE

BIG THING

IN PHYSICS

In 1930, a young physicist named Carl D. Anderson was tasked by his mentor with measuring the energies of cosmic rays—particles arriving at high speed from outer space. Anderson built an improved version of a cloud chamber, a device that visually records the trajectories of particles. In 1932, he saw evidence that confusingly combined the properties of protons and electrons. “A situation began to develop that had its awkward aspects,” he wrote many years after winning a Nobel Prize at the age of 31. Anderson had accidentally discovered antimatter.

Four years after his first discovery, he codiscovered another elementary particle, the muon. This one prompted one physicist to ask, “Who ordered that?”

Over the decades since then, particle physicists have built increasingly sophisticated instruments of exploration. At the apex of these physics-finding machines sits the Large Hadron Collider, which in 2022 started its third operational run. This underground ring, 27 kilometers in circumference and straddling the border between France and Switzerland, was built to slam subatomic particles together at near light speed and test deep theories of the universe. Physicists from around the world turn to the LHC, hoping to find something new. They’re not sure what, but they hope to find it.

It’s the latest manifestation of a rich tradition. Throughout the history of science, new instruments have prompted hunts for the unexpected. Galileo Galilei built telescopes and found Jupiter’s moons. Antonie van Leeuwenhoek built microscopes and noticed “animalcules, very prettily a-moving.” And still today, people peer through lenses and pore through data in search of patterns they hadn’t hypothesized. Nature’s secrets don’t always come with spoilers, and so we gaze into the unknown, ready for anything.

But novel, fundamental aspects of the universe are growing less forthcoming. In a sense, we’ve plucked the lowest-hanging fruit. We know to a good approximation what the building blocks of matter are. The Standard Model of particle physics, which describes the currently known elementary particles, has been in place since the 1970s. Nature can still surprise us, but it typically requires larger or finer instruments, more detailed or expansive data, and faster or more flexible analysis tools.

Those analysis tools include a form of artificial intelligence (AI) called machine learning. Researchers train complex statistical models to find patterns in their data, patterns too subtle for human eyes to see, or too rare for a single human to encounter. At the LHC, which smashes together protons to create immense bursts of energy that decay into other short-lived particles of matter, a theorist might predict some new particle or interaction and describe what its signature would look like in the LHC data, often using a simulation to create synthetic data. Experimentalists would then collect petabytes of measurements and run a machine learning algorithm that compares them with the simulated data, looking for a match. Usually, they come up empty. But maybe new algorithms can peer into corners they haven’t considered.

“An algorithm that can recognize any kind of disturbance would be a win.”

—GEORGIA KARAGIORGI

“You’ve heard probably that there’s a crisis in particle physics,” says Tilman Plehn, a theoretical physicist at Heidelberg University, in Germany. At the LHC and other high-energy physics facilities around the world, the experimental results have failed to yield insights on new physics. “We have a lot of unhappy theorists who thought that their model would have been discovered, and it wasn’t,” Plehn says.

Gregor Kasieczka, a physicist at the University of Hamburg, in Germany, recalls the field’s enthusiasm when the LHC began running in 2008. Back then, he was a young graduate student and expected to see signs of supersymmetry, a theory predicting heavier versions of the known matter particles. The presumption was that “we turn on the LHC, and supersymmetry will jump in your face, and we’ll discover it in the first year or so,” he tells me. Eighteen years later, supersymmetry remains in the theoretical realm. “I think this level of exuberant optimism has somewhat gone.”

The result, Plehn says, is that models for all kinds of things have fallen in the face of data. “And I think we’re going on a different path now.”

That path involves a kind of machine learning called unsupervised learning. In unsupervised learning, you don’t teach the AI to recognize your specific prediction—signs of a particle with this mass and this charge. Instead, you might teach it to find anything out of the ordinary, anything interesting—which could indicate brand new physics. It’s the equivalent of looking with fresh eyes at a starry sky or a slide of pond scum. The problem is, how do you automate the search for something “interesting”?

The Standard Model leaves many questions unanswered. Why do matter particles have the masses they do? Why do neutrinos have mass at all? Where is the particle for transmitting gravity, to match those for

Photo-illustration portraits by Gluekit

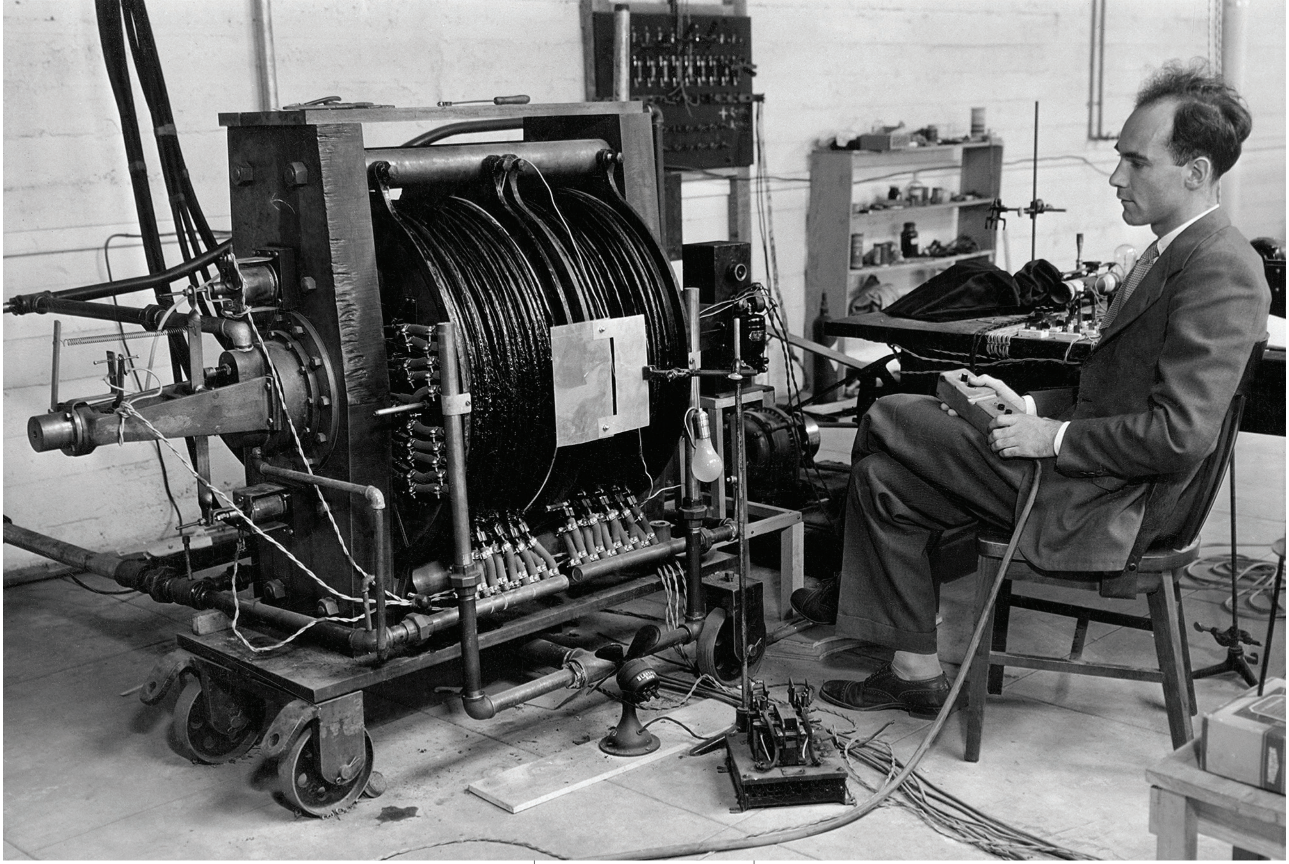

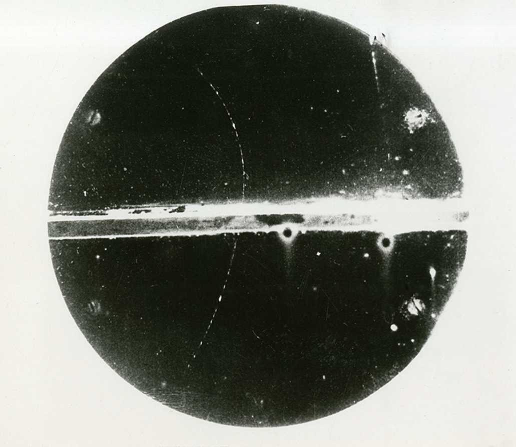

Carl Anderson sits beside the magnet cloud chamber [top] he used to discover the positron. His cloud-chamber photograph [right] from 1932 shows the curved track of a positron, the first known antimatter particle.

CALTECH ARCHIVES & SPECIAL COLLECTIONS

the other forces? Why do we see more matter than antimatter? Are there extra dimensions? What is dark matter—the invisible stuff that makes up most of the universe's matter and that we assume to exist because of its gravitational effect on galaxies? Answering any of these questions could open the door to new physics, or fundamental discoveries beyond the Standard Model.

"Personally, I'm excited for portal models of dark sectors," Kasieczka says, as if reading from a Marvel film script. He asks me to imagine a mirror copy of the Standard Model out there somewhere, sharing only one "portal" particle with the Standard Model we know and love. It's as if this portal particle has a second secret family.

Kasieczka says that in the LHC's third run, scientists are splitting their efforts roughly evenly between measuring more precisely what they know to exist and looking for what they don't know to exist. In some cases, the former could enable the latter. The Standard Model predicts certain particle properties and the relationships between them. For example, it correctly predicted a property of the electron called the magnetic moment to about one part in a trillion. And precise measurements could turn up internal inconsistencies. "Then theorists can say, 'Oh, if I introduce this new particle, it fixes this specific problem that you guys found. And this is how you look for this particle,'" Kasieczka says.

What's more, the Standard Model has occasionally shown signs of cracks. Certain particles containing bottom quarks, for example, seem to decay into other particles in unexpected ratios. Plehn finds the bottom-quark incongruities intriguing. "Year after year, I feel they should go away, and they don't. And nobody has a good explanation," he says. "I wouldn't even know who I would shout at"—the theorists or the experimentalists—"like, 'Sort it out!'"

Exasperation isn't exactly the right word for Plehn's feelings, however. Physicists feel gratified when measurements reasonably agree with expectations, he says. "But I think deep down inside, we always hope that it looks unreasonable. Everybody always looks for the anomalous stuff. Everybody wants to see the standard explanation fail. First, it's fame"—a chance for a Nobel—"but it's also an intellectual challenge, right? You get excited when things don't work in science."

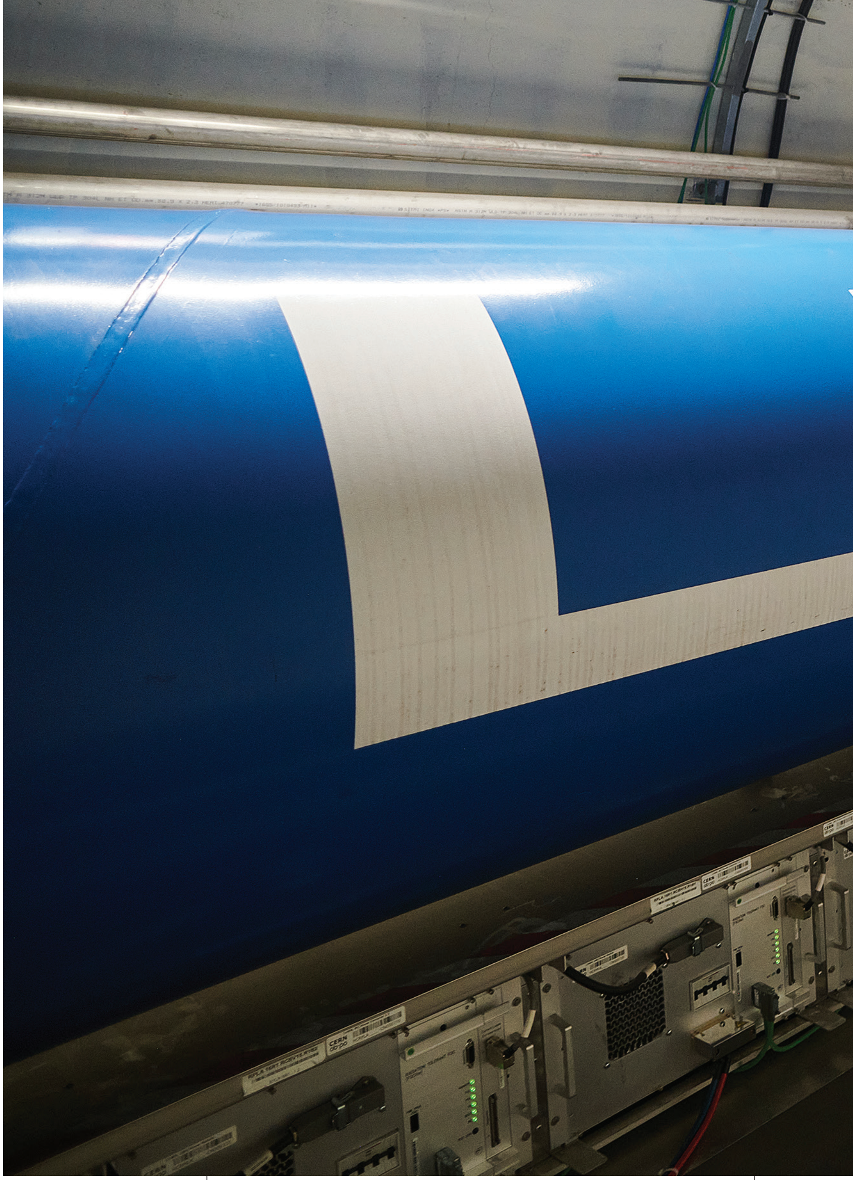

N

Now imagine you had a machine to find all the times things don't work in science, to uncover all the anomalous stuff. That's how researchers are using unsupervised learning. One day over ice cream, Plehn and a friend who works at the software company SAP began discussing autoencoders, one type of unsupervised-learning algorithm. "He tells me that autoencoders are what they use in industry to see if a network was hacked," Plehn remembers. "You have, say, a hundred computers, and they have network traffic. If the network traffic [to one computer] changes all of a sudden, the computer has been hacked, and they take it offline."

Autoencoders are neural networks that start with an input—it could be an image of a cat, or the record of a computer's network traffic—and compress it, like making a tiny JPEG or MP3 file, and then decompress it. Engineers train them to compress and decompress data so that the output matches the input as closely as possible. Eventually a network becomes very good at that task. But if the data includes some items that are relatively rare—such as white tigers, or hacked computers' traffic—the network performs worse on these, because it has less practice with them. The difference between an input and its reconstruction therefore signals how anomalous that input is.

"This friend of mine said, 'You can use exactly our software, right?'" Plehn remembers. "'It's exactly the same question. Replace computers with particles.'" The two imagined feeding the autoencoder signatures of particles from a collider and asking: Are any of these particles not like the others? Plehn continues: "And then we wrote up a joint grant proposal."

It's not a given that AI will find new physics. Even learning what counts as interesting is a daunting hurdle. Beginning in the 1800s, men in lab coats delegated data processing to women, whom they saw as diligent and detail oriented. Women annotated photos of stars, and they acted as "computers." In the 1950s, women were trained to scan bubble chambers, which recorded particle trajectories as lines of tiny bubbles in fluid. Physicists didn't explain to them the theory behind the events, only what to look for based on lists of rules. But, as the Harvard science historian Peter Galison writes in *Image and Logic: A Material Culture of Physics*, his influential account of how physicists' tools shape their discoveries, the task was "subtle, difficult, and anything but routinized," requiring "three-dimensional visual intuition." He goes on: "Even within a single experiment, judgment was required—this was not an algorithmic activity, an assembly line procedure in which action could be specified fully by rules."

Over the last decade, though, one thing we've learned is that AI systems can, in fact, perform tasks once thought to require human intuition, such as mastering the ancient board game Go. So researchers have been testing AI's intuition in physics. In 2019, Kasieczka and his collaborators announced the LHC Olympics 2020, a contest in which participants submitted algorithms to find anomalous events in three sets of (simulated) LHC data. Some teams correctly found the anomalous signal in one dataset, but some falsely reported one in the second set, and they all missed it in the third. In 2020, a research collective called Dark Machines

> "The best scenario could be a new particle. And then a new particle implies a new force."
>
> —JENNIFER NGADIUBA

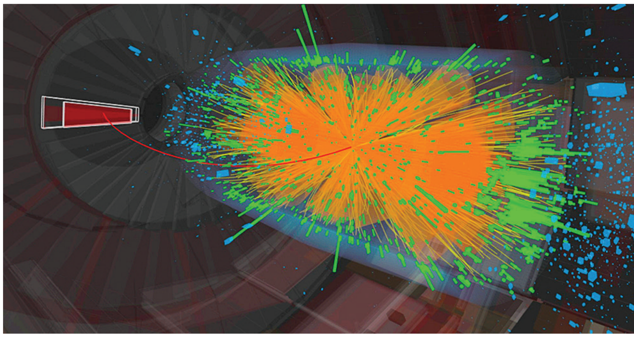

The Large Hadron Collider [left] accelerates protons to near light speed before smashing them together. An image [right] from a collision shows an unusually complex spray of particles, flagged as anomalous by machine learning algorithms.

In the LHC's central data-acquisition room, incoming detector data flows through racks of electronics and field-programmable gate array cards that decide which collision events to keep.

FROM TOP: CERN (2); FERMILAB/CERN

announced a similar competition, which drew more than 1,000 submissions of machine learning models. Decisions about how to score them led to different rankings, showing that there's no best way to explore the unknown.

Another way to test unsupervised learning is to play revisionist history. In 1995, a particle dubbed the top quark turned up at the Tevatron, a particle accelerator at the Fermi National Accelerator Laboratory (Fermilab), in Illinois. But what if it actually hadn't? Researchers applied unsupervised learning to LHC data collected in 2012, pretending they knew almost nothing about the top quark. Sure enough, the AI revealed a set of anomalous events that were clustered together. Combined with a bit of human intuition, they pointed toward something like the top quark.

That exercise underlines the fact that unsupervised learning can't replace physicists just yet. "If your anomaly detector detects some kind of feature, how do you get from that statement to something like a physics interpretation?" Kasieczka says. "The anomaly search is more a scouting-like strategy to get you to look into the right corner.” Georgia Karagiorgi, a physicist at Columbia University, agrees. “Once you find something unexpected, you can’t just call it quits and be like, ‘Oh, I discovered something,’” she says. “You have to come up with a model and then test it.”

Kyle Cranmer, a physicist and data scientist at the University of Wisconsin-Madison who played a key role in the discovery of the Higgs boson particle in 2012, also says that human expertise can’t be dismissed. “There’s an infinite number of ways the data can look different from what you expected,” he says, “and most of them aren’t interesting.” Physicists might be able to recognize whether a deviation suggests some plausible new physical phenomenon, rather than just noise. “But how you try to codify that and make it explicit in some algorithm is much less straightforward,” Cranmer says. Ideally, the guidelines would be general enough to exclude the unimaginable without eliminating the merely unimagined. “That’s gonna be your Goldilocks situation.”

In his 1987 book *How Experiments End*, Harvard’s Galison writes that scientific instruments can “import assumptions built into the apparatus itself.” He tells me about a 1973 experiment that looked for a phenomenon called neutral currents, signaled by an absence of a so-called heavy electron (later renamed the muon). One team initially used a trigger left over from previous experiments, which recorded events only if they produced those heavy electrons—even though neutral currents, by definition, produce none. As a result, for some time the researchers missed the phenomenon and wrongly concluded that it didn’t exist. Galison says that the physicists’ design choice “allowed the discovery of [only] one thing, and it blinded the next generation of people to this new discovery. And that is always a risk when you’re being selective.”

I ask Galison if by automating the search for interesting events, we’re letting the AI take over the science. He rephrases the question: “Have we handed over the keys to the car of science to the machines?” One way to alleviate such concerns, he tells me, is to generate test data to see if an algorithm behaves as expected—as in the LHC Olympics. “Before you take a camera out and photograph the Loch Ness Monster, you want to make sure that it can reproduce a wide variety of colors” and patterns accurately, he says, so you can rely on it to capture whatever comes.

Galison, who is also a physicist, works on the Event Horizon Telescope, which images black holes. For that project, he remembers putting up utterly unexpected test images like Frosty the Snowman so that scientists could probe the system’s general ability to catch something new. “The danger is that you’ve missed out on some crucial test,” he says, “and that the object you’re going to be photographing is so different from your test patterns that you’re unprepared.”

The algorithms that physicists are using to seek new physics are certainly vulnerable to this danger. It helps that unsupervised learning is already being used in many applications. In industry, it’s surfacing anomalous credit-card transactions and hacked networks. In science, it’s identifying earthquake precursors, genome locations where proteins bind, and merging galaxies.

But one difference with particle-physics data is that the anomalies may not be stand-alone objects or events. You’re looking not just for a needle in a haystack; you’re also looking for subtle irregularities in the haystack itself. Maybe a stack contains a few more short stems than you’d expect. Or a pattern reveals itself only when you simultaneously look at the size, shape, color, and texture of stems. Such a pattern might suggest an unacknowledged substance in the soil. In accelerator data, subtle patterns might suggest a hidden force. As Kasieczka and his colleagues write in one paper, “We are not looking for flying elephants, but instead a few extra elephants than usual at the local watering hole.”

Even algorithms that weigh many factors can miss signals—and they can also see spurious ones. The stakes of mistakenly claiming discovery are high. Going back to the hacking scenario, Plehn says, a company might ultimately determine that its network wasn’t hacked; it was just a new employee. The algorithm’s false positive causes little damage. “Whereas if you stand there and get the Nobel Prize, and a year later people say, ‘Well, it was a fluke,’ people would make fun of you for the rest of your life,” he says. In particle physics, he adds, you run the risk of spotting patterns purely by chance in big data, or as a result of malfunctioning equipment.

False alarms have happened before. In 1976, a group at Fermilab led by Leon Lederman, who later won a Nobel for other work, announced the discovery of a particle they tentatively called the Upsilon. The researchers calculated the probability of the signal’s happening by chance as 1 in 50. After further data collection, though, they walked back the discovery, calling the pseudo-particle the Oops-Leon.

“We are not looking for flying elephants but instead a few extra elephants than usual at the local watering hole.”

—GREGOR KASIECZKA

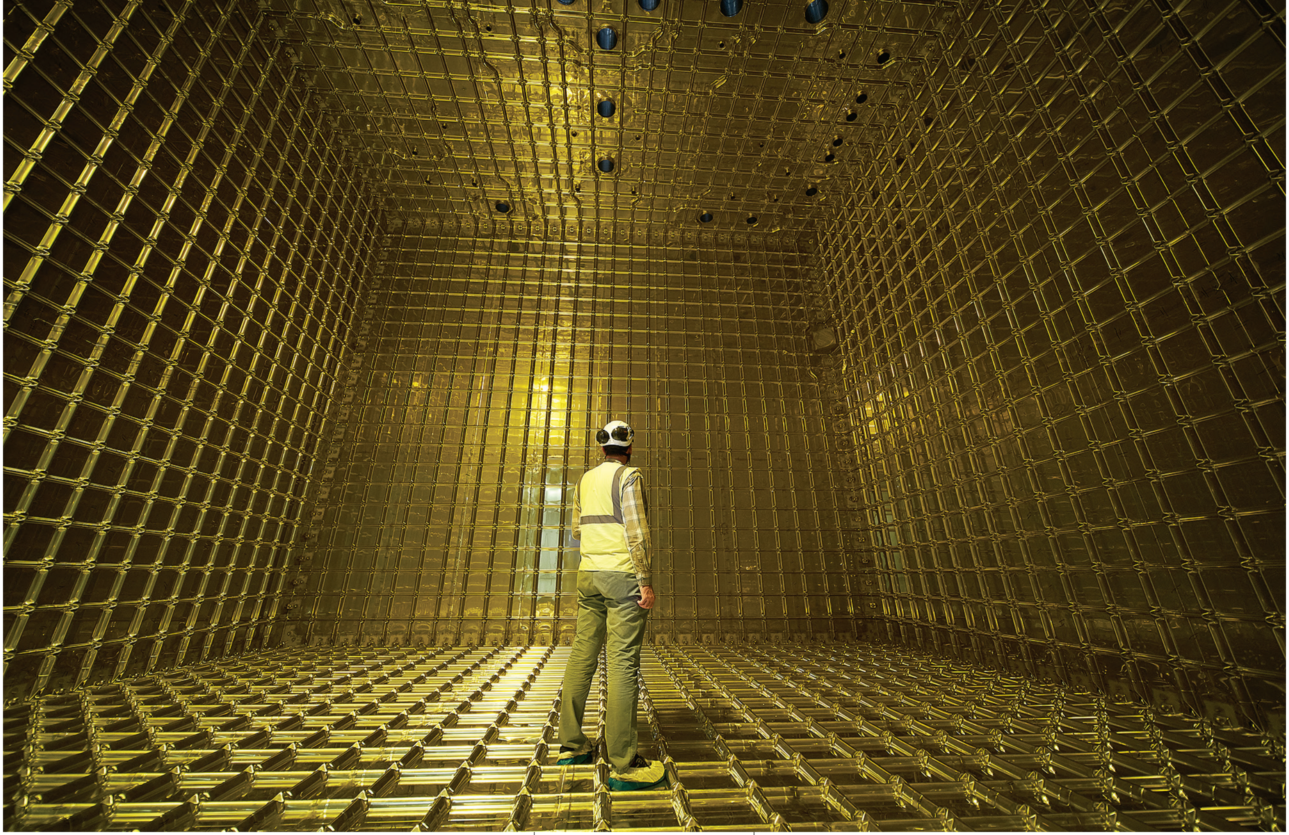

CERN

A researcher stands inside a prototype for the Deep Underground Neutrino Experiment, which is designed to detect rare neutrino interactions.

(Today, particle physicists wait until the chance that a finding is a fluke drops below 1 in 3.5 million, the so-called five-sigma criterion.) And in 2011, researchers at the Oscillation Project with Emulsion-tRacking Apparatus (OPERA) experiment, in Italy, announced evidence for faster-than-light travel of neutrinos. Then, a few months later, they reported that the result was due to a faulty connection in their timing system.

Those cautionary tales linger in the minds of physicists. And yet, even while researchers are wary of false positives from AI, they also see it as a safeguard against them. So far, unsupervised learning has discovered no new physics, despite its use on data from multiple experiments at Fermilab and CERN, home of the LHC. But anomaly detection may have prevented embarrassments like the one at OPERA. “So instead of telling you there’s a new physics particle,” Kasieczka says, “it’s telling you, this sensor is behaving weird today. You should restart it.”

P **Particle physicists** are pushing the limits of not only their computing software but also their computing hardware. The challenge is unparalleled. The LHC produces 40 million particle collisions per second, each of which can produce a megabyte of data. That’s much too much information to store, even if you could save it to disk that quickly. So the two largest detectors each use two-level data filtering. The first layer, called the Level-1 Trigger, or L1T, harvests 100,000 events per second, and the second layer, called the High-Level Trigger, or HLT, plucks 1,000 of those events to save for later analysis. So only one in 40,000 events is potentially seen by human eyes.

HLTs use central processing units (CPUs) like the ones in your desktop computer, running complex machine learning algorithms that analyze collisions based on the number, type, energy, momentum, and angles of the new particles produced. L1Ts, as a first line of defense, must be fast. So the L1Ts rely on integrated circuits called field-programmable gate arrays (FPGAs), which users can reprogram for specialized calculations. The trade-off is that the programming must be relatively simple. The FPGAs can’t easily store and run fancy neural networks; instead they follow scripted rules about, say, what features of a particle collision make it important. In terms of complexity level, it’s the instructions given to the women who scanned bubble chambers, not the women’s brains.

Ekaterina "Katya" Govorkova, a particle physicist at MIT, saw a path toward improving the LHC’s filters, inspired by a board game. Around 2020, she was looking for new physics by comparing precise measurements at the LHC with predictions, using little or no machine learning. Then she watched a documentary about AlphaGo, the program that used machine learning to beat a human Go champion. “For me the moment of realization was when AlphaGo would use some absolutely new type of strategy that humans, who played this game for centuries, hadn’t thought about before,” she says. “So that’s when I thought, we need something like that in physics. We need a genius that can look at the world differently.” New physics may be something we’d never imagine.

Govorkova and her collaborators found a way to compress autoencoders to put them on FPGAs, where they process an event every 80 nanoseconds (less than 10-millionth of a second). (Compression involved pruning some network connections and reducing the precision of some calculations.) They published their methods in *Nature Machine Intelligence* in 2022, and researchers are now using them during the LHC's third run. The new trigger tech is installed in one of the detectors around the LHC's giant ring, and it has found many anomalous events that would otherwise have gone unflagged.

Researchers are currently setting up analysis workflows to decipher why the events were deemed anomalous. Jennifer Ngadiuba, a particle physicist at Fermilab who is also one of the coordinators of the trigger system (and one of Govorkova's coauthors), says that one feature stands out already: Flagged events have lots of jets of new particles shooting out of the collisions. But the scientists still need to explore other factors, like the new particles' energies and their distributions in space. "It's a high-dimensional problem," she says.

Eventually they will share the data openly, allowing others to eyeball the results or to apply new unsupervised-learning algorithms in the hunt for patterns. Javier Duarte, a physicist at the University of California, San Diego, and also a coauthor on the 2022 paper, says, "It's kind of exciting to think about providing this to the community of particle physicists and saying, like, 'Shrug, we don't know what this is. You can take a look.'" Duarte and Ngadiuba note that high-energy physics has traditionally followed a top-down approach to discovery, testing data against well-defined theories. Adding in this new bottom-up search for the unexpected marks a new paradigm. "And also a return of sorts to before the Standard Model was so well established," Duarte adds.

Yet it could be years before we know why AI marked those collisions as anomalous. What conclusions could they support? "In the worst case, it could be some detector noise that we didn't know about," which would still be useful information, Ngadiuba says. "The best scenario could be a new particle. And then a new particle implies a new force."

Duarte says he expects their work with FPGAs to have wider applications. "The data rates and the constraints in high-energy physics are so extreme that people in industry aren't necessarily working on this," he says. "In self-driving cars, usually millisecond latencies are sufficient reaction times. But we're developing algorithms that need to respond in microseconds or less. We're at this technological frontier, and to see how much that can proliferate back to industry will be cool."

"We have a lot of unhappy theorists who thought that their model would have been discovered, and it wasn't." —TILMAN PLEHN

Plehn is also working to put neural networks on FPGAs for triggers, in collaboration with experimentalists, electrical engineers, and other theorists. Encoding the nuances of abstract theories into material hardware is a puzzle. "In this grant proposal, the person I talked to most is the electrical engineer," he says, "because I have to ask the engineer, 'Which of my algorithms fits on your bloody FPGA?'"

Hardware is hard, says Ryan Kastner, an electrical engineer and computer scientist at UC San Diego who works with Duarte on programming FPGAs. What allows the chips to run algorithms so quickly is their flexibility. Instead of programming them in an abstract coding language like Python, engineers configure the underlying circuitry. They map logic gates, route data paths, and synchronize operations by hand. That low-level control also makes the effort "painfully difficult," Kastner says. "It's kind of like you have a lot of rope, and it's very easy to hang yourself."

The next piece of new physics may not pop up at a particle accelerator. It may appear at a detector for neutrinos, particles that are part of the Standard Model but remain deeply mysterious. Neutrinos are tiny, electrically neutral, and so light that no one has yet measured their mass. (The latest attempt, in April, set an upper limit of about a millionth the mass of an electron.) Of all known particles with mass, neutrinos are the universe's most abundant, but also among the most ghostly, rarely deigning to acknowledge the matter around them. Tens of trillions pass through your body every second.

If we listen very closely, though, we may just hear the secrets they have to tell. Karagiorgi, of Columbia, has chosen this path to discovery. Being a physicist is "kind of like playing detective, but where you create your own mysteries," she tells me during my visit to Columbia's Nevis Laboratories, located on a large estate about 20 km north of Manhattan. Physics research began at the site after World War II; one hallway features papers going back to 1951.

Karagiorgi is eagerly awaiting a massive neutrino detector that's currently under construction. Starting in 2028, Fermilab will send neutrinos west through 1,300 km of rock to South Dakota, where they'll occasionally make their existence known in the Deep Underground Neutrino Experiment (DUNE). Why so far away? When neutrinos travel

The Standard Model of Particle Physics

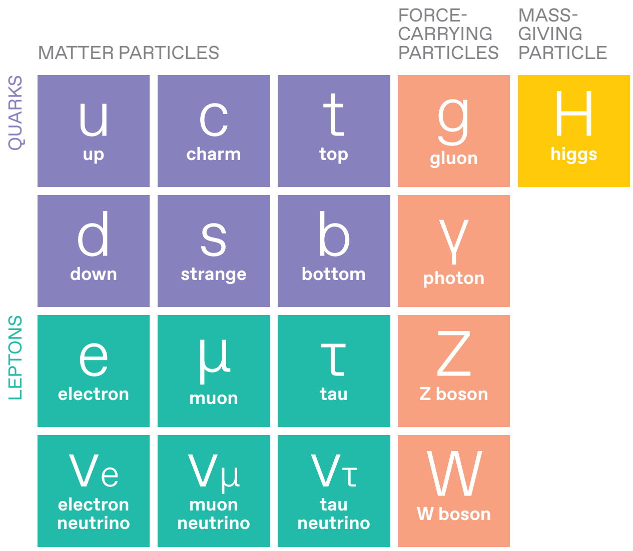

MATTER PARTICLES				FORCE-CARRYING PARTICLES	MASS-GIVING PARTICLE
QUARKS	u up	c charm	t top	g gluon	H higgs
	d down	s strange	b bottom	γ photon	
LEPTONS	e electron	μ muon	τ tau	Z Z boson	
	ν_e electron neutrino	ν_μ muon neutrino	ν_τ tau neutrino	W W boson	

The Standard Model catalogs the known fundamental particles of matter and the forces that govern them, but leaves major mysteries unresolved.

SOURCE: CUSH/WIKIPEDIA

long distances, they have an odd habit of oscillating, transforming from one kind or “flavor” to another. Observing the oscillations of both the neutrinos and their mirror-image antiparticles, antineutrinos, could tell researchers something about the universe’s matter-antimatter asymmetry—which the Standard Model doesn’t explain—and thus, according to the Nevis website, “why we exist.”

“DUNE is the thing that’s been pushing me to develop these real-time AI methods,” Karagiorgi says, “for sifting through the data very, very, very quickly and trying to look for rare signatures of interest within them.” When neutrinos interact with the detector’s 70,000 tonnes of liquid argon, they’ll generate a shower of other particles, creating visual tracks that look like a photo of fireworks.

Even when not bombarding DUNE with neutrinos, researchers will keep collecting data in the off chance that it captures neutrinos from a distant supernova. “This is a massive detector spewing out 5 terabytes of data per second,” Karagiorgi says, “and it’s going to run constantly for a decade.” They will need unsupervised learning to notice signatures that no one was looking for, because there are “lots of different models of how supernova explosions happen, and for all we know, none of them could be the right model for neutrinos,” she says. “To train your algorithm on such uncertain grounds is less than ideal. So an algorithm that can recognize any kind of disturbance would be a win.”

Deciding in real time which 1 percent of 1 percent of data to keep will require FPGAs. Karagiorgi’s team is preparing to use them for DUNE, and she walks me to a computer lab where they program the circuits. In the FPGA lab, we look at nondescript circuit boards sitting on a table. “So what we’re proposing is a scheme where you can have something like a hundred of these boards for DUNE deep underground that receive the image data frame by frame,” she says. This system could tell researchers whether a given frame resembled TV static, fireworks, or something in between.

Neutrino experiments, like many particle-physics studies, are very visual. When Karagiorgi was a postdoc, automated image processing at neutrino detectors was still in its infancy, so she and collaborators would often resort to visual scanning (bubble-chamber style) to measure particle tracks. She still asks undergrads to hand-scan as an educational exercise. “I think it’s wrong to just send them to write a machine learning algorithm. Unless you can actually visualize the data, you don’t really gain a sense of what you’re looking for,” she says. “I think it also helps with creativity to be able to visualize the different types of interactions that are happening, and see what’s normal and what’s not normal.”

Back in Karagiorgi’s office, a bulletin board displays images from *The Cognitive Art of Feynman Diagrams*, an exhibit for which the designer Edward Tufte created wire sculptures of the physicist Richard Feynman’s schematics of particle interactions. “It’s funny, you know,” she says. “They look like they’re just scribbles, right? But actually, they encode quantitatively predictive behavior in nature.” Later, Karagiorgi and I spend a good 10 minutes discussing whether a computer or a human could find Waldo without knowing what Waldo looked like. We also touch on the 1964 Supreme Court case in which Justice Potter Stewart famously declined to define obscenity, saying “I know it when I see it.” I ask whether it seems weird to hand over to a machine the task of deciding what’s visually interesting. “There are a lot of trust issues,” she says with a laugh.

On the drive back to Manhattan, we discuss the history of scientific discovery. “I think it’s part of human nature to try to make sense of an orderly world around you,” Karagiorgi says. “And then you just automatically pick out the oddities. Some people obsess about the oddities more than others, and then try to understand them.”

Reflecting on the Standard Model, she called it “beautiful and elegant,” with “amazing predictive power.” Yet she finds it both limited and limiting, blinding us to colors we don’t yet see. “Sometimes it’s both a blessing and a curse that we’ve managed to develop such a successful theory.” ■

CERN's Walter Wuensch [top left] says the work on FLASH radiation therapy to treat cancer is "generating a lot of excitement" at the particle physics research facility. Hardware [top right and bottom left] developed to boost electrons to ultrahigh power over short distances is leading to compact linear accelerators suitable for a clinical setting, like Theryq's FLASHKNiFE system [bottom right], which is being trialed to treat superficial tumors.

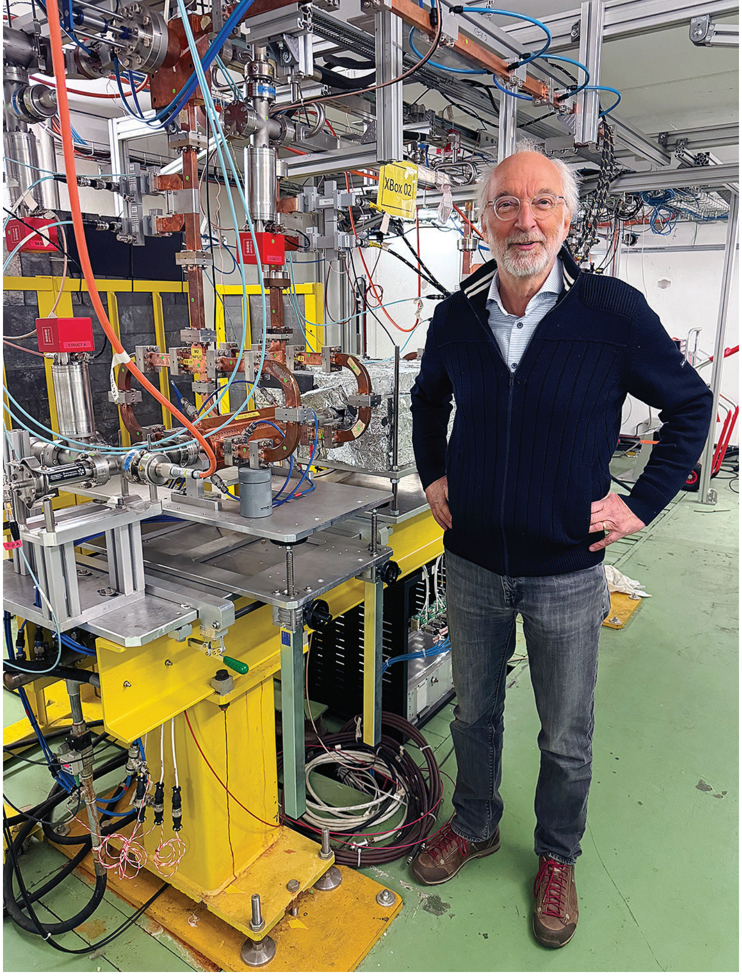

The Millisecond That Could Change Cancer Treatment

FLASH THERAPY USES A SINGLE BLAST OF ULTRAHIGH-POWER RADIATION TO KILL TUMORS WHILE SPARING HEALTHY TISSUE

By Tom Clynes

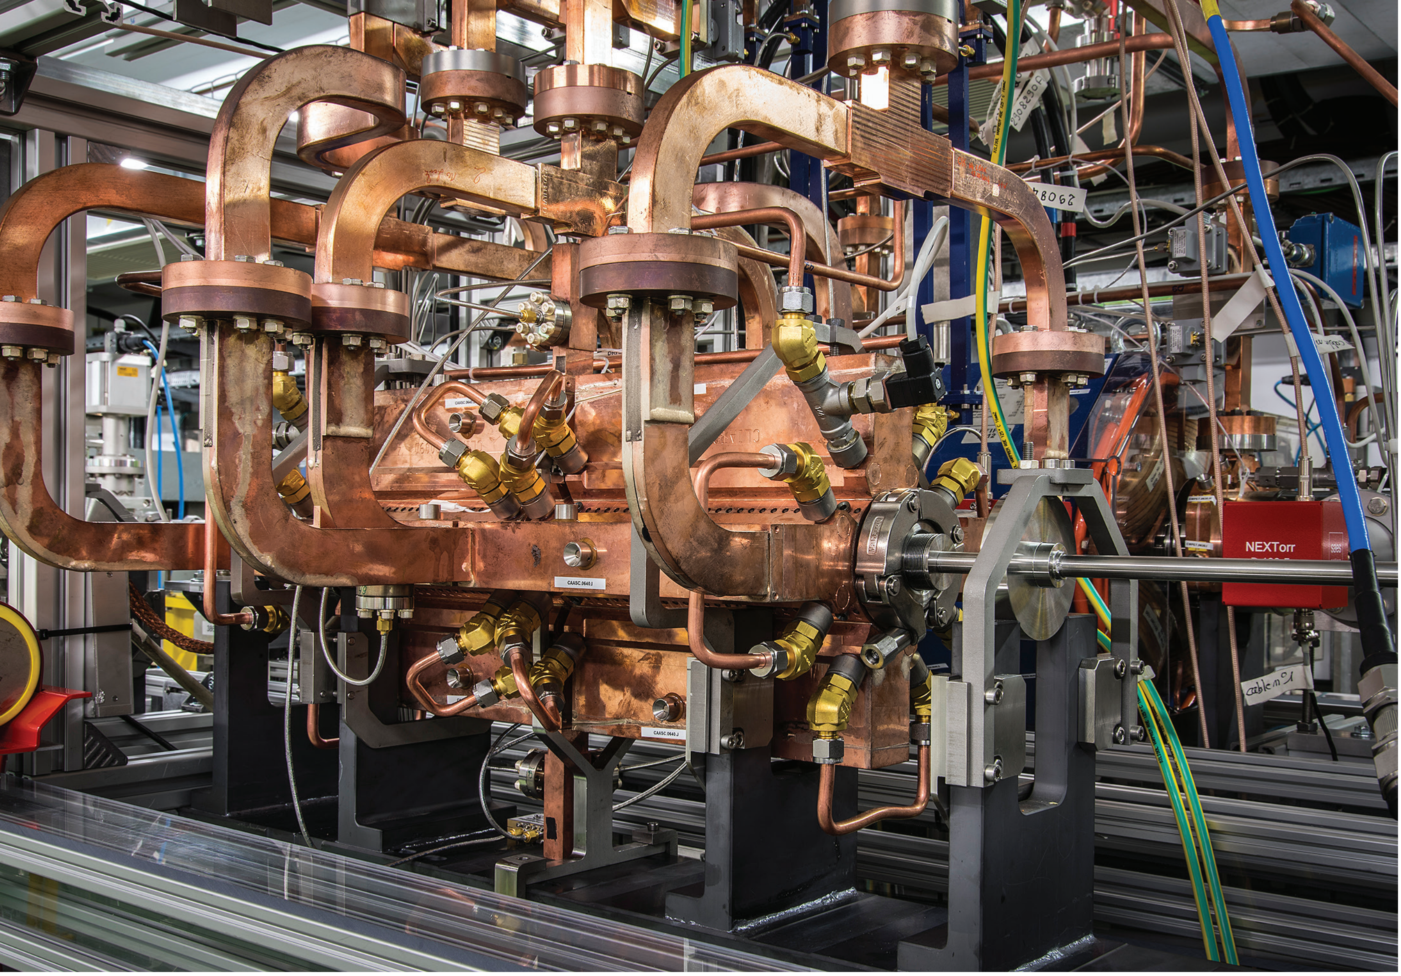

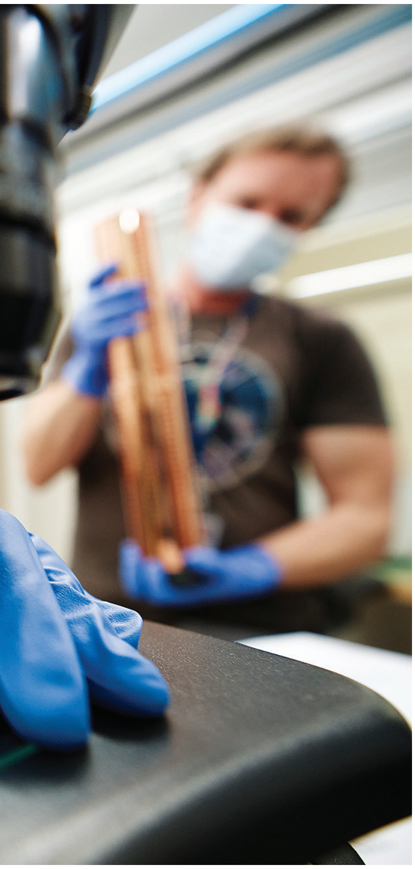

CLOCKWISE FROM BOTTOM LEFT: CERN (3); THERYQ

Inside a cavernous hall at the Swiss-French border,

the air hums with high voltage and possibility. From his perch on the wraparound observation deck, physicist Walter Wuensch surveys a multimillion-dollar array of accelerating cavities, klystrons, modulators, and pulse compressors—hardware being readied to drive a new generation of linear particle accelerators.

Wuensch has spent decades working with these machines to crack the deepest mysteries of the universe. Now he and his colleagues are aiming at a new target: cancer. Here at CERN (the European Organization for Nuclear Research) and other particle-physics labs, scientists and engineers are applying the tools of fundamental physics to develop a technique called FLASH radiotherapy that offers a radical and counterintuitive vision for treating the disease.

Radiation therapy has been a cornerstone of cancer treatment since shortly after Wilhelm Conrad Röntgen discovered X-rays in 1895. Today, more than half of all cancer patients receive it as part of their care, typically in relatively low doses of X-rays delivered over dozens of sessions. Although this approach often kills the tumor, it also wreaks havoc on nearby healthy tissue. Even with modern precision targeting, the potential for collateral damage limits how much radiation doctors can safely deliver.

FLASH radiotherapy flips the conventional approach on its head, delivering a single dose of ultrahigh-power radiation in a burst that typically lasts less than one-tenth of a second. In study after study, this technique causes significantly less injury to normal tissue than conventional radiation does, without compromising its antitumor effect.

At CERN, which I visited last July, the approach is being tested and refined on accelerators that were never intended for medicine. If ongoing experiments here and around the world continue to bear out results, FLASH could transform radiotherapy—delivering stronger treatments, fewer side effects, and broader access to lifesaving care.

"It's generating a lot of excitement," says Wuensch, a researcher at CERN's Linear Electron Accelerator for Research (CLEAR) facility. "We accelerator people are thinking, Oh, wow, here's an application of our technology that has a societal impact which is more immediate than most high-energy physics."

The breakthrough that led to FLASH emerged from a line of experiments that began in the 1990s at Institut Curie in Orsay, near Paris. Researcher Vincent Favaudon was using a low-energy electron accelerator to study radiation chemistry. Targeting the accelerator at mouse lungs, Favaudon expected the radiation to produce scar tissue, or fibrosis. But when he exposed the lungs to ultrafast blasts of radiation, at doses a thousand times as high as what's used in conventional radiation therapy, the expected fibrosis never appeared.

Puzzled, Favaudon turned to Marie-Catherine Vozenin, a radiation biologist at Curie who specialized in radiation-induced fibrosis. "When I looked at the slides, there was indeed no fibrosis, which was very, very surprising for this type of dose," recalls Vozenin, who now works at Geneva University Hospitals, in Switzerland.

The pair expanded the experiments to include cancerous tumors. The results upended a long-held trade-off of radiotherapy: the idea that you can't destroy a tumor without also damaging the host. "This differential effect is really what we want in radiation oncology, not damaging normal tissue but killing the tumors," Vozenin says.

They repeated the protocol across different types of tissue and tumors. By 2014, they had gathered enough evidence to publish their findings in *Science Translational Medicine*. Their experiments confirmed that delivering an ultrahigh dose of 10 grays or more in less than a tenth of a second could eradicate tumors in mice while leaving surrounding healthy tissue virtually unharmed. For comparison, a typical chest X-ray delivers about 0.1 milligray, while a session of conventional radiation therapy might deliver a total of about 2 grays per day. (The authors called the effect "FLASH" because of the quick, high doses involved, but it's not an acronym.)

Many cancer experts were skeptical. The FLASH effect seemed almost too good to be true. "It didn't get a lot of

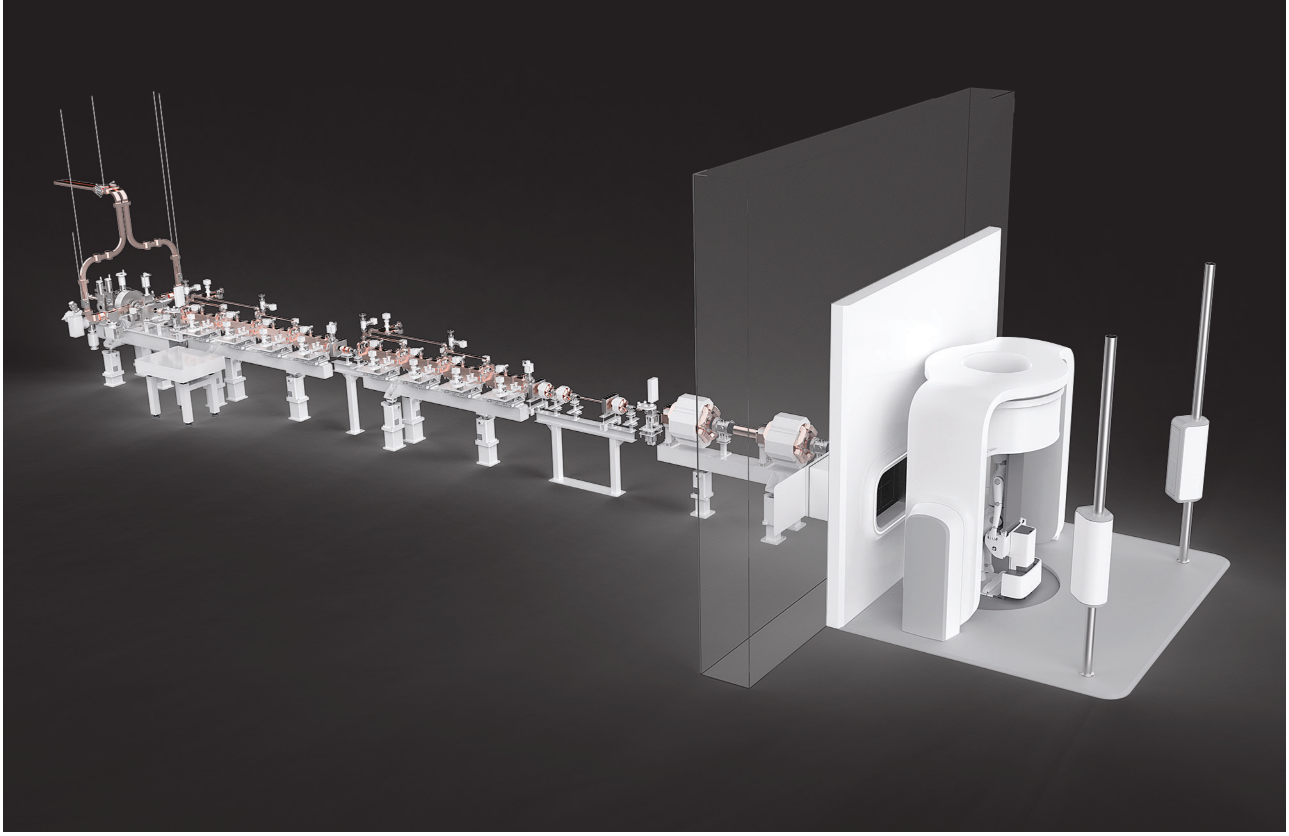

Theryq's FLASHDEEP system, under development with CERN and the company's clinical partners, has a 13.5-meter-long, 140-MeV linear accelerator. That's strong enough to treat tumors at depths of up to about 20 centimeters in the body. The patient will remain in a supported standing position during the split-second irradiation.

traction at first," recalls Billy Loo, a Stanford radiation oncologist specializing in lung cancer. "They described a phenomenon that ran counter to decades of established radiobiology dogma."

But in the years since then, researchers have observed the effect across a wide range of tumor types and animals—beyond mice to zebra fish, fruit flies, and even a few human subjects, with the same protective effect in the brain, lungs, skin, muscle, heart, and bone.

Why this happens remains a mystery. "We have investigated a lot of hypotheses, and all of them have been wrong," says Vozenin. Currently, the most plausible theory emerging from her team's research points to metabolism: Healthy and cancerous cells may process reactive oxygen species—unstable oxygen-containing molecules generated during radiation—in very different ways.

At the time of the first FLASH publication, Loo and his team at Stanford were also focused on dramatically speeding up radiation delivery. But Loo wasn't chasing a radiobiological breakthrough. He was trying to solve a different problem: motion.

"The tumors that we treat are always moving targets," he says. "That's particularly true in the lung, where because of breathing motion, the tumors are constantly moving."

To bring FLASH therapy out of the lab and into clinical use, researchers like Vozenin and Loo needed machines capable of delivering fast, high doses with pinpoint precision deep inside the body. Most early studies relied on lower-energy electron beams like Favaudon's 4.5-megaelectron-volt Kinetron—sufficient for surface tumors, but unable to reach more than a few centimeters into a human body. Treating deep-seated cancers in the lung, brain, or abdomen would require far higher particle energies.

They also needed an alternative to conventional X-rays. In a clinical linac, X-ray photons are produced by dumping high-energy electrons into a bremsstrahlung target—an inherently inefficient process that wastes most of the beam power as heat and makes it extremely difficult to reach the ultrahigh dose rates required for FLASH. High-energy electrons, by contrast, can be switched on and off within milliseconds. And because they have a charge and can be steered by magnets, electrons can be precisely guided to reach tumors deep within the body. (Researchers are also investigating protons and carbon ions; see the sidebar "What's the Best Particle for FLASH Therapy?")

Loo turned to the SLAC National Accelerator Laboratory in Menlo Park, Calif., where physicist Sami Gamal-Eldin Tantawi was redefining how electromagnetic waves move through linear accelerators. Tantawi's findings allowed scientists to precisely control how energy is delivered to particles—paving the way for compact, efficient, and finely tunable machines. It was exactly the kind of technology FLASH therapy would need to target tumors deep inside the body.

Meanwhile, Vozenin and other European researchers turned to CERN, best known for its 27-kilometer Large Hadron Collider (LHC) and the 2012 discovery of the Higgs boson, the "God particle," which gives other particles their mass. CERN is also home to a range of smaller linear accelerators—including CLEAR, where Wuensch and his team are adapting high-energy physics tools for medicine.

THERYQ

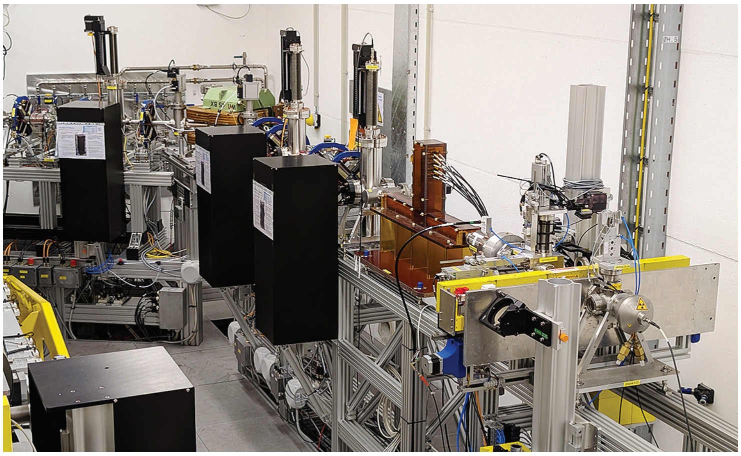

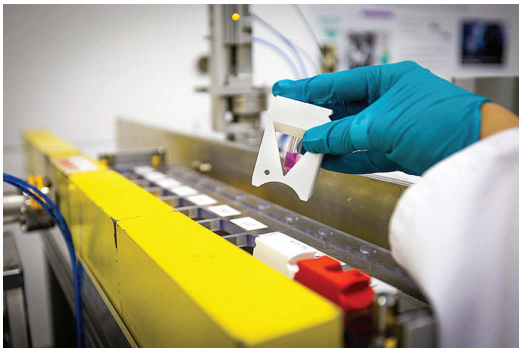

At Germany's Photo Injector Test Facility in Zeuthen (PITZ), the electron-beam accelerator [left] is used to irradiate the biological targets in early-stage animal tests of FLASH radiotherapy [right].

Unlike the LHC, which loops particles around a massive ring to build up energy before smashing them together, linear accelerators like CLEAR send particles along a straight, one-time path. That setup allows for greater precision and compactness, making it ideal for applications like FLASH.

At the heart of the CLEAR facility, Wuensch points out the 200-MeV linear accelerator with its 20-meter beamline. This is "a playground of creativity," he says, for the physicists and engineers who arrive from all over the world to run experiments.

The process begins when a laser pulse hits a photocathode, releasing a burst of electrons that form the initial beam. These electrons travel through a series of precisely machined copper cavities, where high-frequency microwaves push them forward. The electrons then move through a network of magnets, monitors, and focusing elements that shape and steer them toward the experimental target with submillimeter precision.

Instead of a continuous stream, the electron beam is divided into nanosecond-long bunches—billions of electrons riding the radio-frequency field like surfers. Inside the accelerator's cavities, the field flips polarity 12 billion times per second, so timing is everything: Only electrons that arrive perfectly in phase with the accelerating wave will gain energy. That process repeats through a chain of cavities, each giving the bunches another push, until the beam reaches its final energy of 200 MeV.

Much of this architecture draws directly from the Compact Linear Collider study, a decades-long CERN project aimed at building a next-generation collider. The proposed CLIC machine would stretch 11 kilometers and collide electrons and positrons at 380 gigaelectron volts. To do that in a linear configuration—without the multiple passes around a ring like the LHC—CERN engineers have had to push for extremely high acceleration gradients to boost the electrons to high energies over relatively short distances—up to 100 megavolts per meter.

Wuensch leads me to a large experimental hall housing prototype structures from the CLIC effort, and points out the microwave devices that now help drive FLASH research. Though the future of CLIC as a collider remains uncertain, its infrastructure is already yielding dividends: smaller, high-gradient accelerators that may one day be as suited for curing cancer as they are for smashing particles.

The power behind the high gradients comes from CERN's Xboxes, the X-band RF systems that dominate the experimental hall. Each Xbox houses a klystron, modulator, pulse

How to Measure Radiation Doses

RADIATION THERAPY uses a variety of units to refer to the amount of energy received by the patient. Here are the main ones under the International System of Units, or SI.

Gray (Gy): A measure of the absorbed dose—that is, how much radiation energy is absorbed by the body. One gray equals 1 joule of radiation energy per kilogram of matter. FLASH delivers a single dose of 40 Gy or more in a fraction of a second. Conventional radiation therapy, by contrast, may deliver a total dose of 40 to 80 Gy but over the course of several weeks.

Sievert (Sv): A measure of the effective dose—that is, the health effects of the radiation, with different types of ionizing radiation (gamma rays, X-rays, alpha particles, and so on) having different effects. One sievert equals 1 joule per kilogram weighted for the biological effectiveness of the radiation and the tissues exposed.

LEFT: FRIEDER MUELLER; RIGHT: MWFK

compressor, and waveguide network to generate and shape the microwave pulses. The pulse compressors store energy in resonant cavities and then release it in a microsecond burst, producing peaks of up to 200 megawatts; if it were continuous, that's enough to power at least 40,000 homes. The Xboxes let researchers fine-tune the power, timing, and pulse shape.

According to Wuensch, many of the recent accelerator developments were enabled by advances in computer simulation and high-precision three-dimensional machining. These tools allow the team to iterate quickly, designing new accelerator components and improving beam control with each generation.

Still, real-world challenges remain. The power demands are formidable, as are the space requirements; for all the talk of its "compact" design, the original CLIC was meant to span kilometers. Obviously, a hospital needs something that's actually compact.

"A big challenge of the project," says Wuensch, "is to transform this kind of technology and these kinds of components into something that you can imagine installing in a hospital, and it will run every day reliably."

To that end, CERN researchers have teamed up with Lausanne University Hospital (known by its French acronym, CHUV) and the French medical technology company Theryq to design a hospital facility capable of treating large and deep-seated tumors with the very short time scales needed for FLASH and scaled down to fit in a clinical setting.

Theryq's research center and factory are located in southern France, near the base of Montagne Sainte-Victoire, a jagged spine of limestone that Paul Cézanne painted dozens of times, capturing its shifting light and form.

"The solution that we are trying to develop here is something which is extremely versatile," says Ludovic Le Meunier, CEO of the growing company. "The ultimate goal is to be able to treat any solid tumor anywhere in the body, which is about 90 percent of the cancer these days."

Theryq's push to bring FLASH radiotherapy from the lab to the clinic has followed a three-pronged rollout, with each device engineered for a specific depth and clinical use. The first machine, FLASHKNiFE, was unveiled in 2020. Designed for superficial tumors and intraoperative use, the system delivers electron beams at 6 or 9 MeV. A prototype installed that same year at CHUV is conducting a phase-two trial for patients with localized skin cancer.

More recently, Theryq launched FLASHLAB, a compact, 7-MeV platform for radiobiology research.

The company's most ambitious system, FLASHDEEP, is still under development with CERN and Theryq's clinical partners. The 13.5-meter-long electron source will deliver very high-energy electrons of as much as 140 MeV up to 20 centimeters inside the body in less than 100 milliseconds.

What's the Best Particle for FLASH Therapy?

EVEN AS RESEARCH on FLASH radiotherapy advances, a central question remains: What kind of particle will deliver it best? The main contenders are electrons, protons, and carbon ions. Each has distinct implications for cost, complexity, and clinical reach.

Electrons—long used to treat surface tumors and to generate X-rays—are light, nimble particles, far easier to control than protons or carbon ions. At low energies, they stop quickly in tissue, but new high-energy systems can drive electrons deeper. Now researchers are working on machines that combine multiple high-energy beams at different angles to let doctors sculpt radiation doses that match the tumor's shape.

That principle underpins Billy Loo's PHASER (Pluridirectional High-energy Agile Scanning Electron Radiotherapy) system, developed at Stanford and SLAC and licensed to startup TibaRay. An array of high-efficiency linacs generates X-ray beams from many directions at once. Their high output overcomes the inefficiency of electron-to-photon conversion to deliver the dose at FLASH speed. Beam convergence at the tumor and electronic shaping conform the dose in three dimensions, producing uniform coverage with relatively simple infrastructure.

Protons have led the way in early clinical trials, largely because existing proton therapy centers can be adapted to deliver FLASH doses. In 2020, UC Health, in Cincinnati, launched the first human FLASH trial to use proton beams, to treat cancer that had metastasized to bones. "If I want to be pragmatic, the proton beam is ready to go, so let's move with what we have," says Geneva University Hospitals' Marie-Catherine Vozenin.

Protons can penetrate up to 30 centimeters, reaching deep-seated tumors. But the delivery of protons in a continuous beam limits the dose rates. Also, proton systems are far larger and more expensive than, say, X-ray machines, which will likely constrain their availability to specialized centers.

Carbon ions, used in a handful of elite facilities, offer even higher precision and biological effectiveness compared to electrons and protons. Their Bragg peak—a sudden deposition of energy at a specific depth—makes them appealing for deep or complex tumors. But that unmatched precision comes at a steep price, with each facility costing upward of US $300 million. —T.C.

An integrated CT scanner, built into the patient-positioning system developed by Leo Cancer Care, captures images that stream directly into the treatment-planning software, enabling precise calculation of the radiation dose. “Before we actually trigger the beam or the treatment, we make stereo images to verify at the very last second that the tumor is exactly where it should be,” says Theryq technical manager Philippe Liger.

While CERN’s CLEAR accelerator has been instrumental in characterizing FLASH parameters, researchers seeking to study FLASH in living organisms must look elsewhere: CERN doesn’t allow animal experiments on-site. That’s one reason why a growing number of scientists are turning to PITZ, the Photo Injector Test Facility in Zeuthen, a leafy lakeside suburb of Berlin.

PITZ is part of Germany’s national accelerator lab and is responsible for developing the electron source for the European X-ray Free-Electron Laser. Now PITZ is emerging as a hub for FLASH research, with an unusually tunable accelerator and a dedicated biomedical lab to ensure controlled conditions for preclinical studies.

“The biggest advantage of our facility is that we can do a very stepwise, very defined and systematic study of dose rates,” says Anna Grebinyk, a biochemist who heads the new biomedical lab, “and systematically optimize the FLASH effect to see where it gets the best properties.”

The experiments begin with zebra-fish embryos, prized for early-stage studies because they’re transparent and develop rapidly. After the embryos, researchers test the most promising parameters in mice. To do that, the PITZ team uses a small-animal radiation research platform, complete with CT imaging and a robotic positioning system adapted from CERN’s CLEAR facility.

What sets PITZ apart is the flexibility of its beamline. The 30-meter accelerator system steers electrons with micrometer precision, producing electron bunches with exceptional brightness and emittance—a metric of beam quality. “We can dial in any distribution of bunches we want,” says Frank Stephan, group leader at PITZ. “That gives us tremendous control over time structure.”

Timing matters. At PITZ, the laser-struck photocathode generates electron bunches that are accelerated immediately, at up to 60 million volts per meter. A fast electromagnetic kicker system acts as a high-speed gatekeeper, selectively deflecting individual electron bunches from a high-repetition beam and steering them according to researchers’ needs. This precise, bunch-by-bunch control is essential for fine-tuning beam properties for FLASH experiments and other radiation therapy studies.

“The idea is to make the complete treatment within one millisecond,” says Stephan. “But of course, you have to [trust] that within this millisecond, everything works fine. There is not a chance to stop [during] this millisecond. It has to work.”

Regulating the dose remains a major technical hurdle. The ionization chambers used in standard radiotherapy can’t respond accurately when dose rates spike hundreds of times higher in a matter of microseconds. So researchers are developing new detector systems to precisely measure these bursts and keep pace with the extreme speed of FLASH delivery.

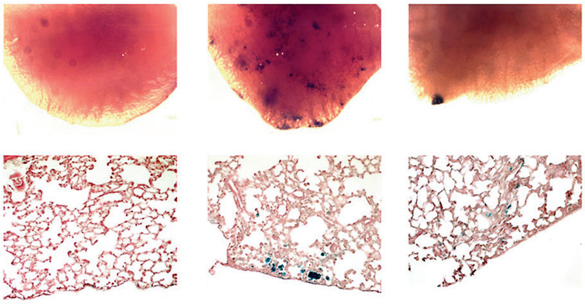

Although many cancer experts were skeptical about the FLASH effect on healthy tissue when it was first announced in 2014, numerous studies have since confirmed and expanded on those results. In a 2020 paper, a lung tissue sample taken four months after being exposed to conventional radiotherapy [center] shows many more dark spots indicating scarring than a sample exposed to FLASH [right]. The nonirradiated sample [left] is the control.

VINCENT FAVAUDON/AMERICAN ASSOCIATION FOR CANCER RESEARCH

Beyond its therapeutic potential, FLASH may also open new windows to illuminate cancer biology. “What is really, really superinteresting, in my opinion,” says Vozenin, “is that we can use FLASH as a tool to understand the difference between normal tissue and tumors. There must be something we’re not aware of that really distinguishes the two—and FLASH can help us find it.” Identifying those differences, she says, could lead to entirely new interventions, not just with radiation, but also with drugs.

Vozenin’s team is currently testing a hypothesis involving long-lived proteins present in healthy tissue but absent in tumors. If those proteins prove to be key, she says, “we’re going to find a way to manipulate them—and perhaps reverse the phenomenon, even [turn] a tumor back into a normal tissue.”

Proponents of FLASH believe it could help close the cancer-care gap worldwide; in low-income countries, only about 10 percent of patients have access to radiotherapy, and in middle-income countries, only about 60 percent of patients do, according to the International Atomic Energy Agency. Because FLASH treatment can often be delivered in a single brief session, it could spare patients from traveling long distances for weeks of treatment and allow clinics to treat many more people.

High-income countries stand to benefit as well. Fewer sessions mean lower costs, less strain on radiotherapy facilities, and fewer side effects and disruptions for patients.

The big question now is, How long will it take? Researchers I spoke with estimate that FLASH could become a routine clinical option in about 10 years—after the completion of remaining preclinical studies and multiphase human trials, and as machines become more compact, affordable, and efficient. Much of the momentum comes from a growing field of startups competing to build devices, but the broader scientific community remains remarkably open and collaborative.

“Everyone has a relative who knows about cancer because of their own experience,” says Stephan. “My mother died of it. In the end, we want to do something good for mankind. That’s why people work together.” ■

ATLAS

Atlas is an advanced humanoid robot built for fully autonomous work in real-world industrial environments. While earlier versions relied on hydraulic actuators, the latest model is entirely electric, enabling faster, quieter, and more-precise motion. Its highly articulated body allows it to bend, twist, and move in ways that go beyond human physical limits.

CREATOR
Boston Dynamics

YEAR
2024

COUNTRY
United States

CATEGORY
Humanoids

FEATURES
Agility, AI, dexterity, full autonomy, vision

PHOTO: BOSTON DYNAMICS

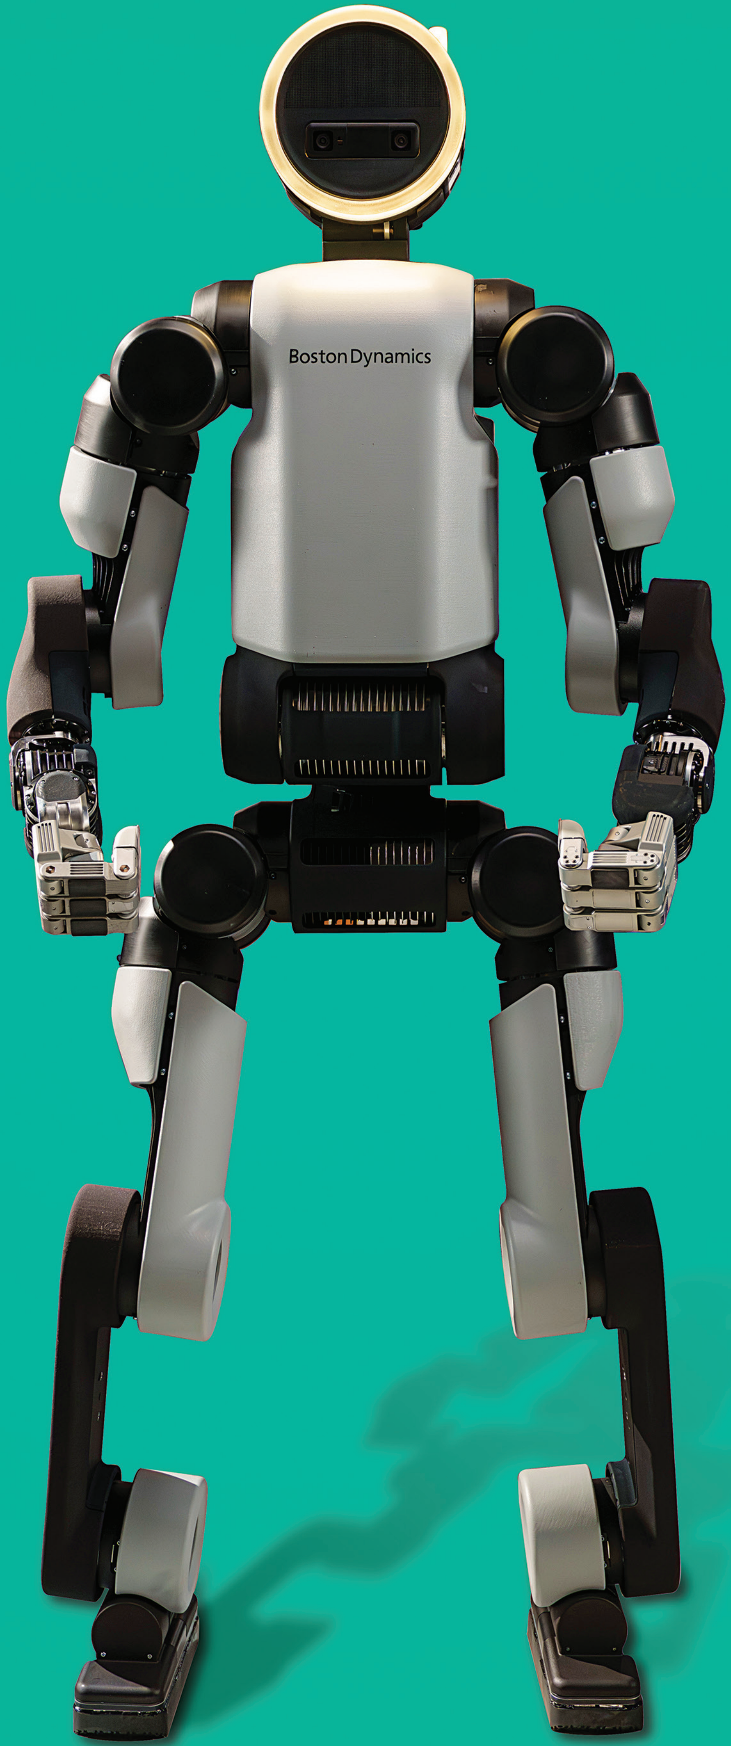

To learn more about Atlas and explore over 270 real-world robots, with thousands of photos, videos, and interactive animations, scan the QR code or visit the IEEE Robots Guide at robotsguide.com

ROBOTS

ENERGY

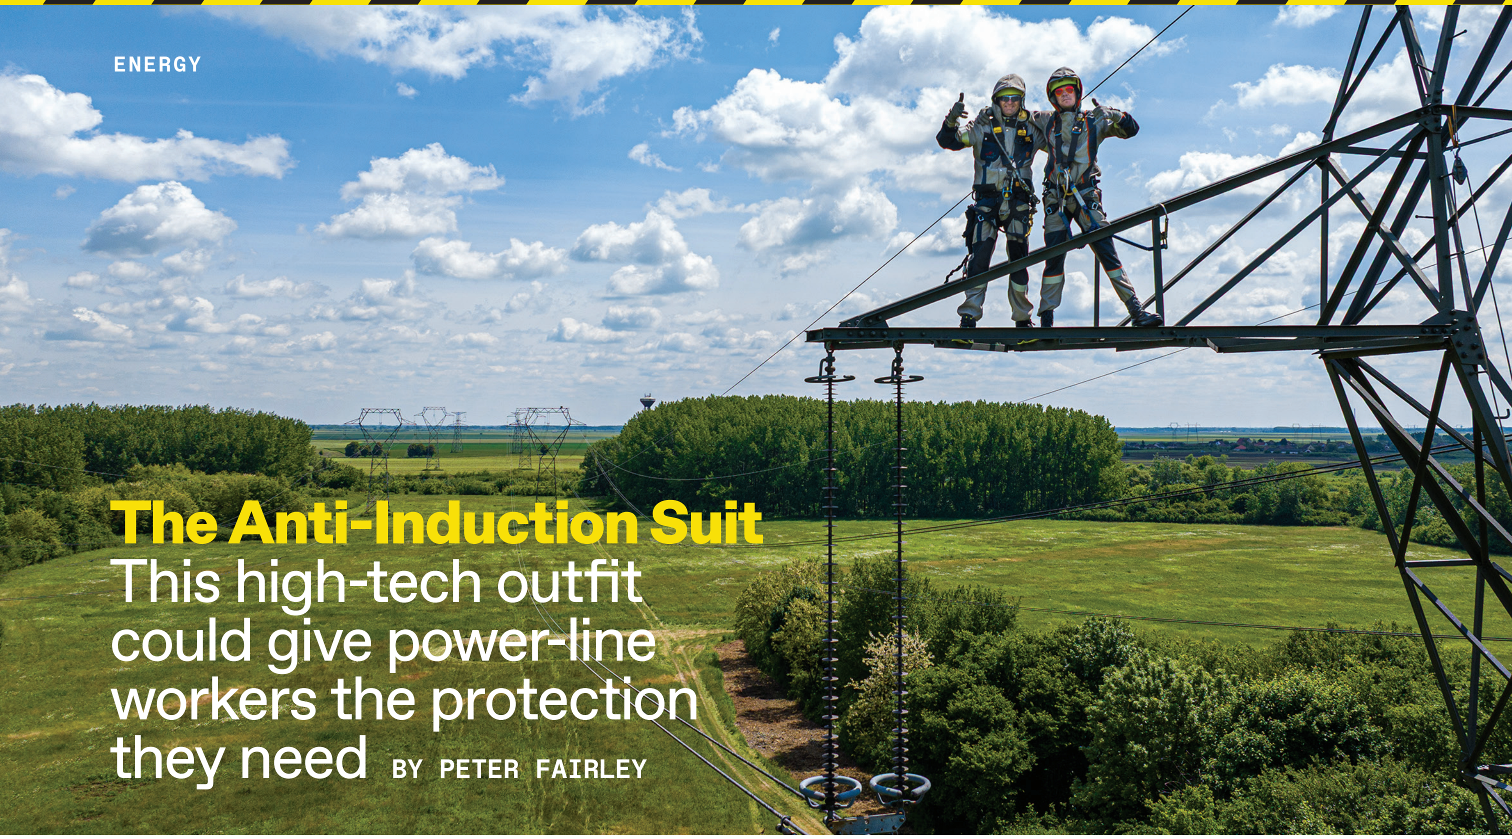

The Anti-Induction Suit

This high-tech outfit could give power-line workers the protection they need

BY PETER FAIRLEY

IN 2018, JUSTIN KROPP was working on a transmission circuit in Southern California when disaster struck. Grid operators had earlier shut down the 115-kilovolt circuit, but six high-voltage lines that shared the corridor were still operating, and some of their power sneaked onto the de-energized wires he was working on. That rogue current shot to the ground through Kropp's body and his elevated work platform, killing the 32-year-old father of two.

"It went in both of his hands and came out his stomach where he was leaning against the platform rail," says Justin's father, Barry Kropp, who is himself a retired line worker. "Justin got hung up on the wire. When they finally got him on the ground, it was too late."

Justin's accident was caused by induction: a hazard that occurs when an electric or magnetic field causes current to flow through equipment whose intended power supply has been cut off. Safety practices seek to prevent such induction shocks by grounding all conductive objects in a work zone, giving electricity alternative paths. But accidents happen. In Justin's case, his platform unexpectedly swung into the line before it could be grounded.

Adding a layer of defense against induction injuries is the motivation behind Budapest-based Electrostatics' specialized conductive jumpsuits, which are designed to protect against burns, cardiac fibrillation, and other ills. "If my boy had been wearing one, I know he'd be alive today," says the elder Kropp, who purchased a line-worker safety training business after Justin's death. The Mesa, Ariz.–based company, Electrical Safety Consulting International (ESCI), now distributes those suits.

Eduardo Ramirez Bettoni, one of the developers of the suits, dug into induction risk after a series of major accidents in the United States in 2017 and 2018, including Justin Kropp's. At the time, he was principal engineer for transmission and substation standards at Minneapolis-based Xcel Energy. In talking to Xcel line workers and fellow safety engineers, he sensed that the accident cluster might be the tip of an iceberg. And when he and two industry colleagues scoured data from the U.S. Bureau of Labor Statistics, they found 81 induction accidents between 1985 and 2021 and 60 deaths, which they documented in a 2022 report.

"Unfortunately, it is really common. I would say there are hundreds of induction contacts every year in the United States alone," says Ramirez Bettoni, who is now technical director of R&D for Houston-based Powell Industries, a power-distribution-equipment firm. He bets that such "contacts"—exposures to dangerous levels of induction—are increasing as grid operators boost grid capacity by squeezing additional circuits into transmission corridors.

Electrostatics' suits are an enhancement of the standard protective gear that line workers wear when their tasks involve working close to or even touching energized live lines, or "bare-hands" work. Both are interwoven with conductive materials such as stainless steel threads, which form a Faraday cage that shields the wearer against the lines' electric fields. But the standard suits have limited capacity to shunt current because usually they don't need to. Like a bird on a wire, bare-hands workers are electrically floating, rather than grounded, so current largely bypasses them via the line itself.

Backed by a US $250,000 investment from Xcel in 2019, Electrostatics adapted

ABOVE: ELECTROSTATICS; RIGHT: BME HVL (4)

← Linemen working on a high-voltage transmission tower wear Electrostatics' suits to protect against grounding failures. But the dangers of induced current aren't limited to workers; it can also ignite dry vegetation and unleash wildfires.

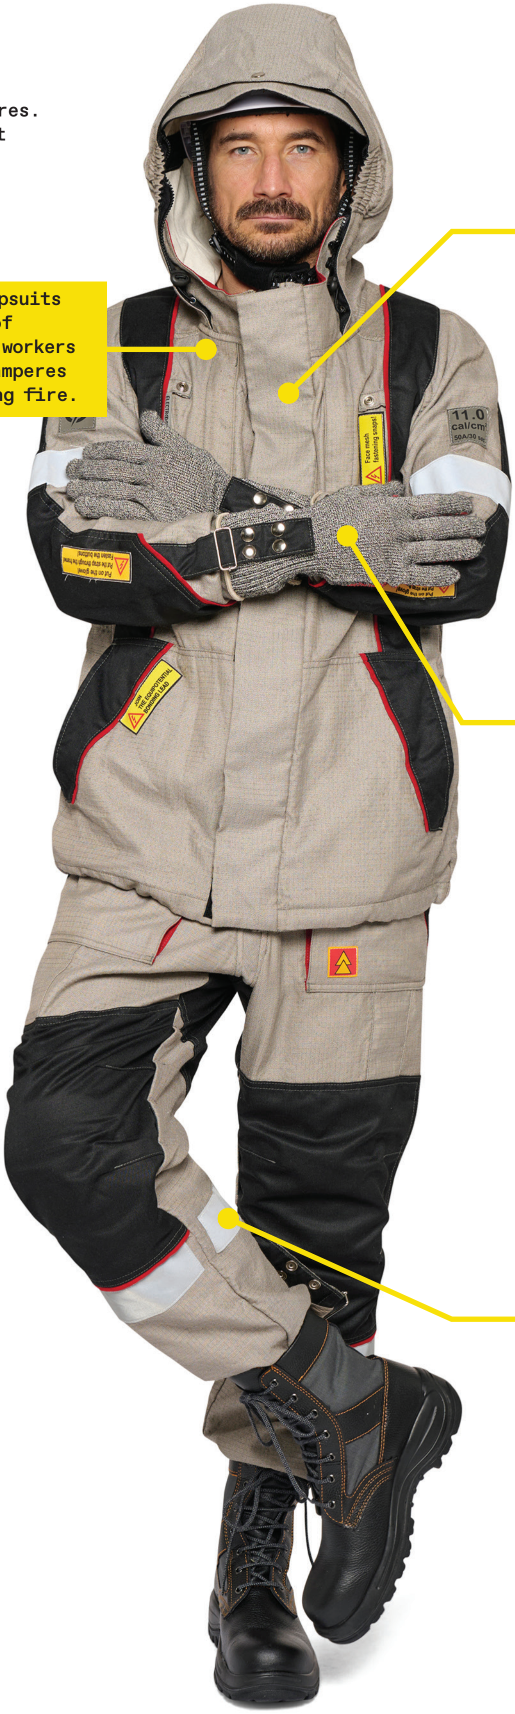

Specialized conductive jumpsuits made by Electrostatics, of Budapest, can protect line workers by absorbing hundreds of amperes of current without catching fire.

its standard suits by adding low-resistance conductive straps that pass current around a worker's body. "When I'm touching a conductor with one hand and the other hand is grounded, the current will flow through the straps to get out," says Bálint Németh, Electrostatics' CEO and director of the High Voltage Laboratory at Budapest University of Technology and Economics.

The company began selling the suits in 2023 and they have since been adopted by over a dozen transmission operators in the United States and Europe, as well as other countries including Canada, Indonesia, and Turkey. They cost $5,214 in the United States.

Electrostatics' suits had to meet a crucial design threshold: keeping body exposure below the 6-milliampere "let-go" threshold, beyond which electrocuted workers become unable to remove themselves from a circuit. "If you lose control of your muscles, you're going to hold onto the conductor until you pass out or possibly die," says Ramirez Bettoni.

The gear, which includes the suit, gloves, and socks, protects against 100 amperes for 10 seconds and 50 A for 30 seconds. It also has insulation to protect against heat created by high current and flame retardants to protect against electric arcs.

Kropp, Németh, and Ramirez Bettoni, are hoping that developing industry standards for induction safety gear, including ones published in October, will broaden their use. Meanwhile, the recently enacted Justin Kropp Safety Act in California, for which the elder Kropp lobbied, mandates automated defibrillators at power-line work sites. ■

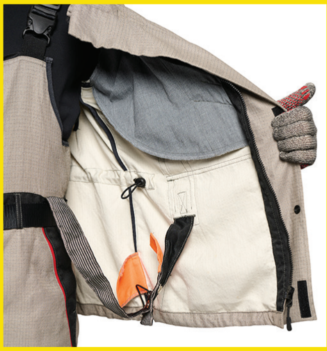

A strapping system links all the elements of the suit—the jacket, trousers, gloves, and socks—and guides current through a controlled path outside the body.

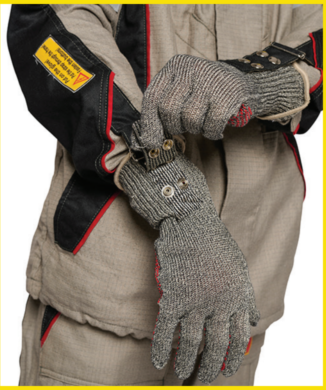

Conductive gloves connect to the sleeves of the suit via a strapping system that provides a controlled bypass path for unexpected current.

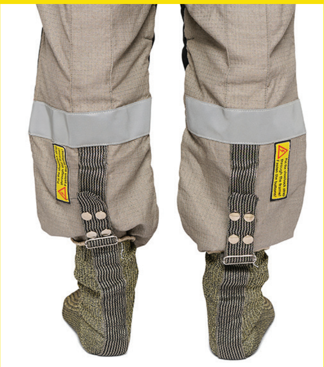

Conductive socks that are connected to the trousers complete the protective suit.

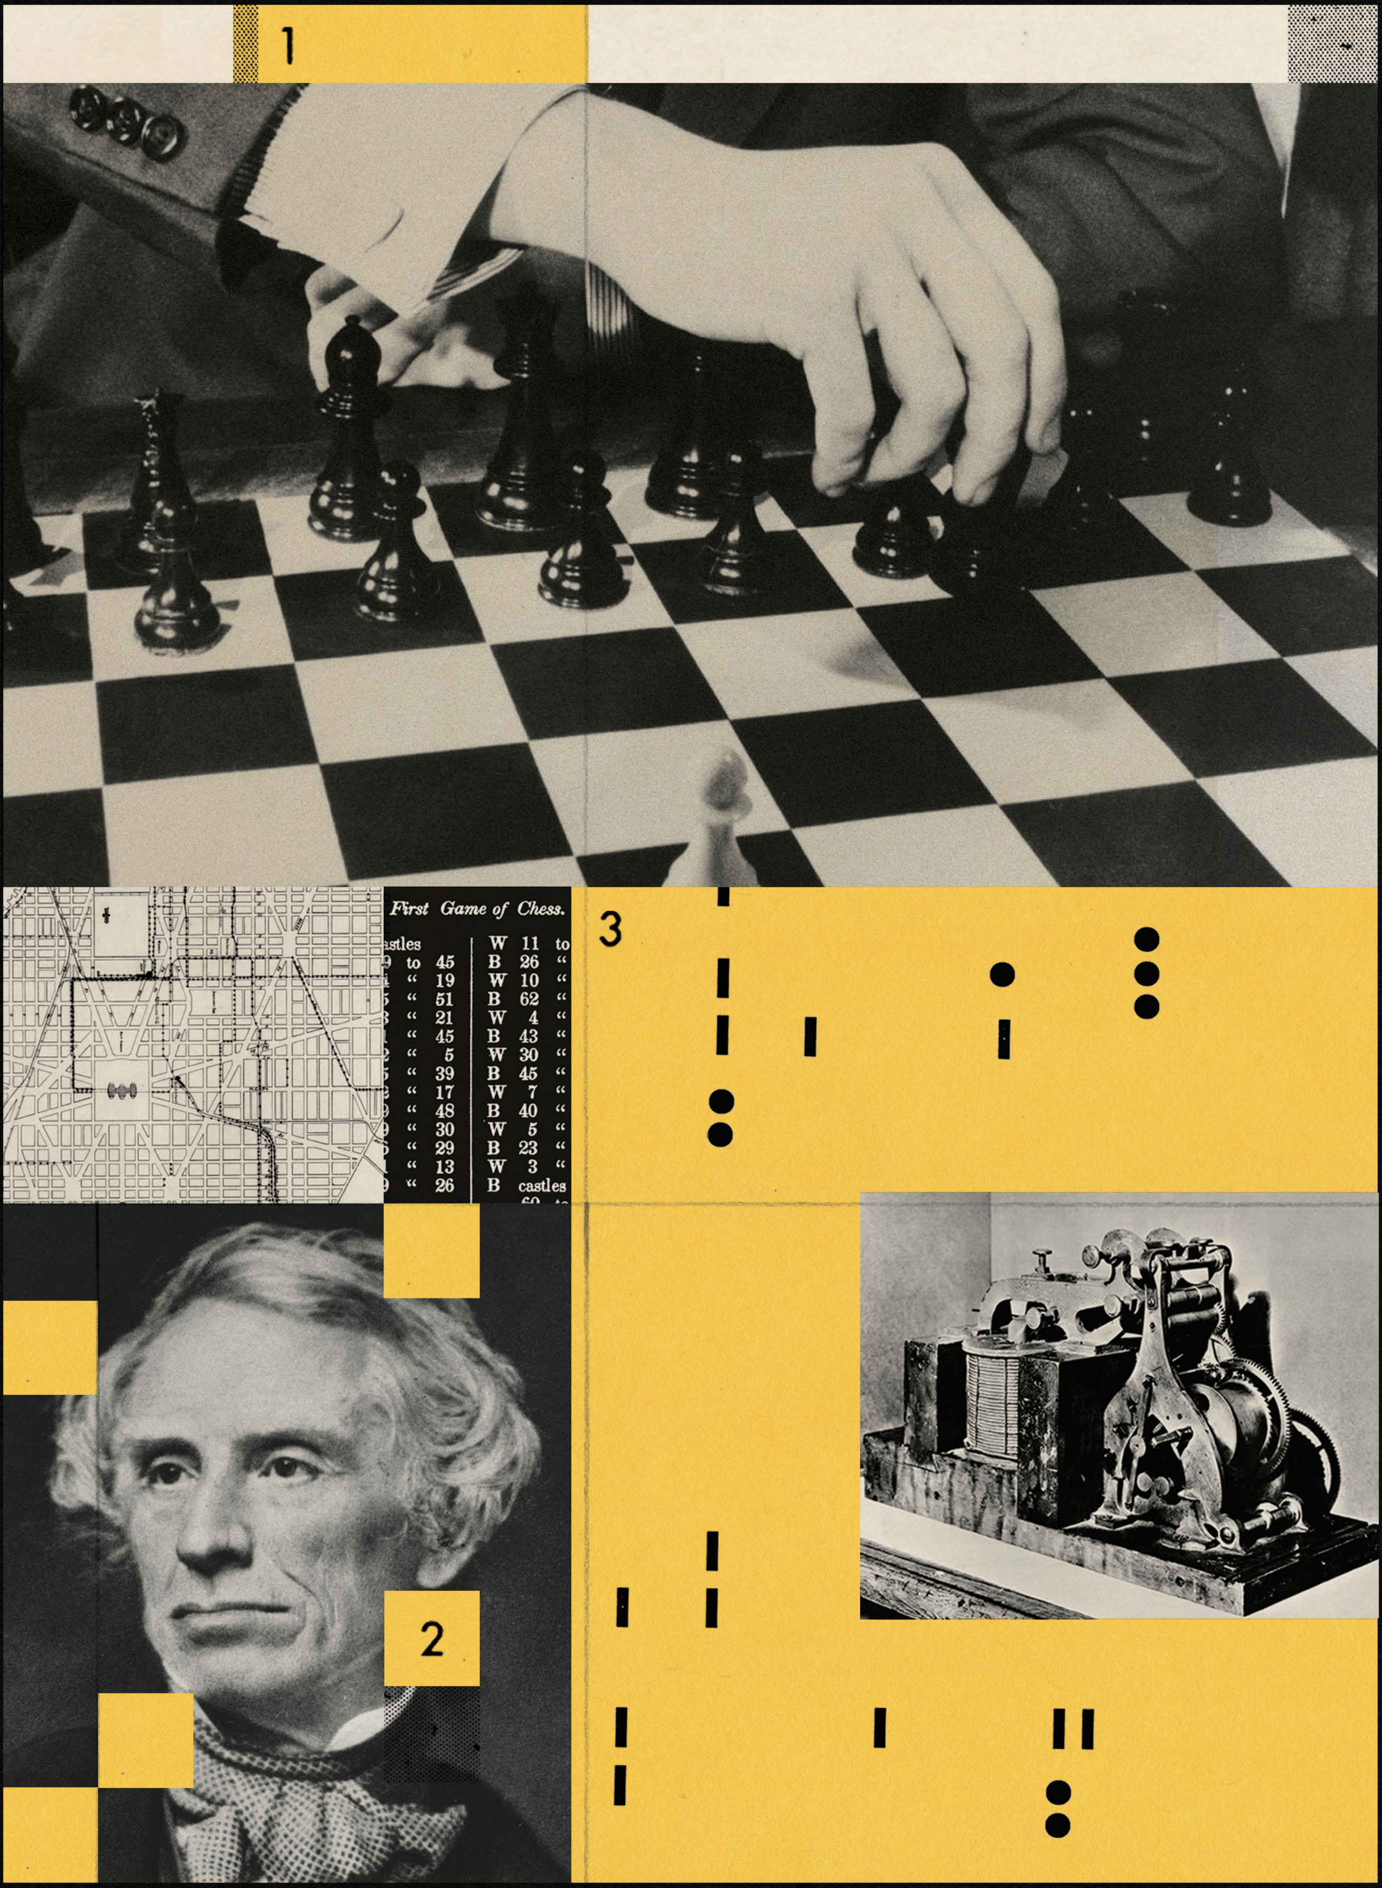
1
First Game of Chess.
astles W 11 to
9 to 45 B 26 "
" 19 W 10 "
" 51 B 62 "
" 21 W 4 "
" 45 B 43 "
" 5 W 30 "
" 39 B 45 "
" 17 W 7 "
" 48 B 40 "
" 30 W 5 "
" 29 B 23 "
" 13 W 3 "
" 26 B castles
3
2

IN 1844 CHESS WAS ALREADY ONLINE

THE FIRST AMERICAN TELEGRAPH LINE WAS ONLY SIX MONTHS OLD WHEN IT BECAME A VIRTUAL CHESSBOARD

By Greg Uyeno | Illustration by Mike McQuade

ON 18 NOVEMBER 1844, the Washington Chess Club challenged its counterparts in Baltimore to a match. Two teams were organized, and at 4 p.m. on 26 November, the first game commenced with three consulting members to a side. Washington began conventionally, pushing a pawn to the center of the board. Baltimore immediately responded by mirroring the move. But this was unlike any chess game ever played before. The Baltimoreans were still in Baltimore, the Washingtonians were still in Washington, D.C, 60 kilometers away, and they were playing by electrical telegraph.

Successive moves were transmitted over the new Baltimore-Washington telegraph line, the first in the United States. Samuel Morse and company had inaugurated the line in May of that year with the message "What hath God wrought."

One chess game led to another, and play continued on and off for days. Records of the games are incomplete and sometimes inconsistent—181 years later, it's unclear who exactly dreamt up chess over wire and why. But thanks in part to historical documents at the Smithsonian Institution, we know enough about the people involved and the operation of the early telegraph to have a sense of the proceedings. We know that Morse would cite chess in lobbying Congress to fund the extension of the telegraphic network to New York via Philadelphia. And we know that there was much more chess by telegraph to come.

Not simply a novelty or a one-off tech demo, telegraph chess eventually became a well-known and sometimes joked-about trend in the United States and Britain, writes historian Simone Müller-Pohl. Chess by telegraph also prefigured chess played through other means of telecommunications. There are records of recreational and serious games played over radio, telephone lines, and satellite, and through online forums, email, apps, and dedicated live services. Most recently, chess has evolved into an esport. Last year, chess joined the likes of *Call of Duty*, *Street Fighter*, and *Rocket League* at the 2025 Esports World Cup in Riyadh, Saudi Arabia.

The number of adults worldwide who play chess regularly is often estimated at around 600 million, and many of them use whatever means available to play games across long distances with friends, rivals, and strangers. Indeed, the 1,500-year-old game and the latest in telecommunications always seem to find each other, starting just months after the first telegraph was built in the United States, when chess went electric.

The Baltimore-Washington telegraph was financed in 1843 with US $30,000 (about $1.3 million today) appropriated by Congress, with the help of Morse's business partner, Francis O.J. Smith, who had supported the project in 1838 while still a sitting congressman from Maine. By late 1844, a bill to extend the line to New York was before the U.S. House of Representatives. In at least one way, drawing the attention of legislators to the new line was easy—the Washington end moved back and forth between the Capitol building

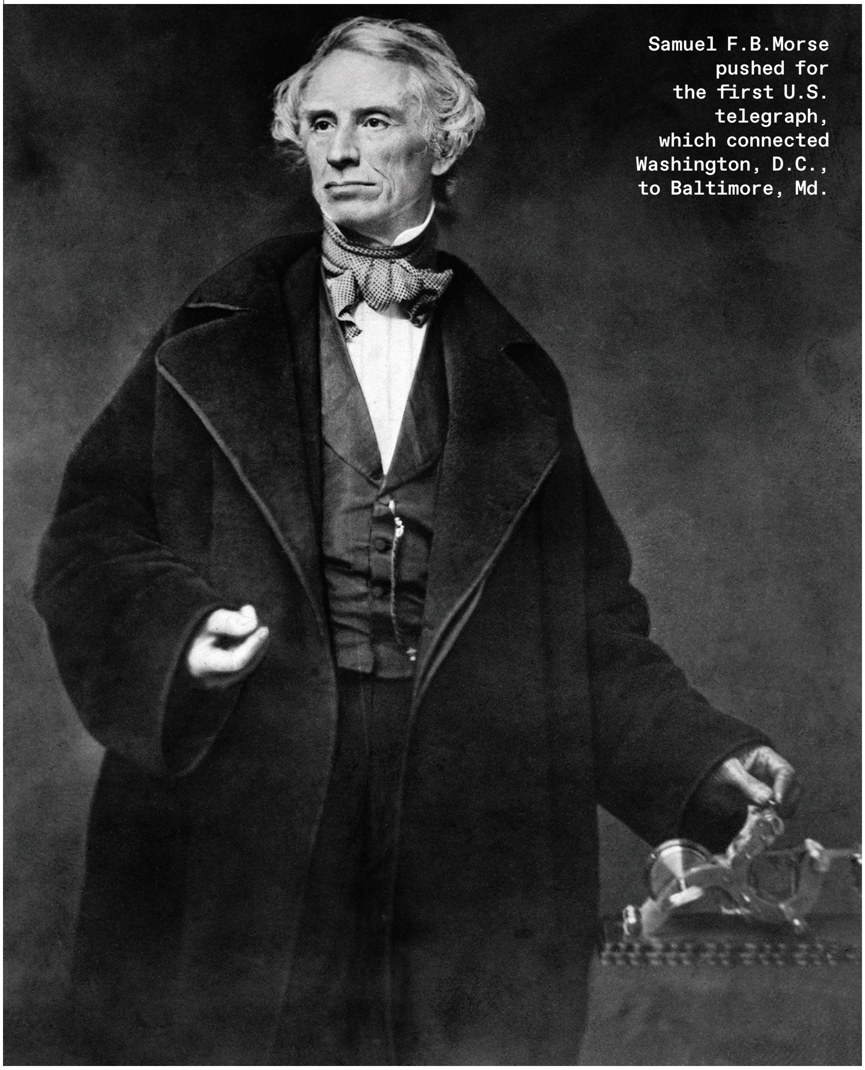
Samuel F.B.Morse pushed for the first U.S. telegraph, which connected Washington, D.C., to Baltimore, Md.

MATHEW B. BRADY/LIBRARY OF CONGRESS

Alfred Vail [right], Morse's associate, was instrumental in organizing the first telegraph chess match and kept detailed notes [below] on messages sent over the line.

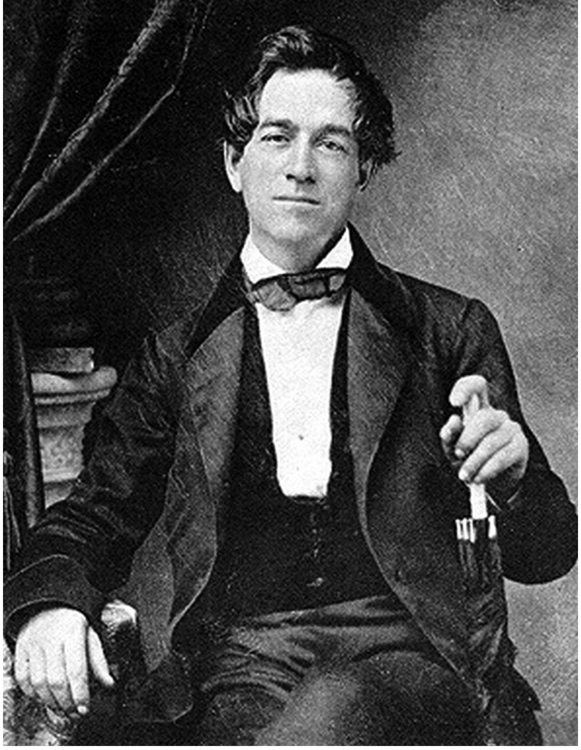

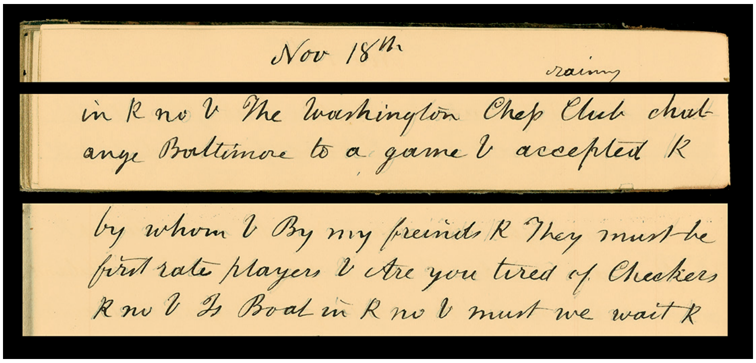

Nov 18th
rainy

in R no V The Washington Chess Club chal
ange Baltimore to a game V accepted R

by whom V By my freinds R They must be
first rate players V Are you tired of Checkers
R no V Is Board in R no V must we wait R

LEFT: SMITHSONIAN INSTITUTION ARCHIVES;
RIGHT: ZOOM HISTORICAL/ALAMY

and the post office, near the present-day National Portrait Gallery. If you were a lawmaker in Washington at the time, the telegraph would've been hard to miss.

But perhaps they needed more persuading. Orrin S. Wood, a telegraph operator, seems to have thought so. On 5 December 1844, Wood wrote a letter to his brother-in-law, engineer Ezra Cornell, who had worked on the line and would go on to cofound Western Union:

"We have had considerable excitement in playing chess between this place and Baltimore for the last 2 or 3 weeks.... I am inclined to think that Congress will do something for Prof Morse as very many of them appear to be very much interested with [chess]."

A week later, Morse wrote to George M. Bibb, Secretary of the Treasury, to lobby for the funding. The telegraph could relay congressional news, presidential convention results, and the whereabouts of wanted criminals, he argued. He also played the chess card:

"To show the variety of the operations of the telegraph, a game of draughts [checkers], and several games of chess, have been played between the cities of Baltimore and Washington, with the same ease as if the players were seated at the same table."

Chess had even been played on rainy nights, he noted. The telegraph's continued operation in both inclement weather and darkness compared favorably with optical telegraphy. Such systems, popular in France, consisted of regularly spaced flag towers that relayed messages by semaphore; they were costly to build and operate and only worked in daylight and good weather.

While Morse appealed to people's interest in telegraphic chess, the game itself didn't obviously begin with promotional intent. It began with checkers. We know this because Morse's associate Alfred Vail kept a "Journal of the Magnetic Telegraph between Washington and Baltimore" (now part of the Vail Telegraph Collection at the Smithsonian) in which he meticulously recorded messages sent over the wire for posterity.

On 15 November 1844, Vail in Washington instructed Henry J. Rogers in Baltimore to "get a checkerboard and let us play a game tomorrow morning." Vail promised to send instructions by regular mail on the five o'clock train. Rogers came up with the idea of using numbered squares to communicate locations on the board. Later that day, Rogers announced that John Wills, a journalist with the *Baltimore Patriot*, would play in his place.

The next morning, before the checkers game began, Vail recorded a telegraphic exchange between himself and Rogers, in which Vail suggests the game is for private enjoyment, and he would prefer that Wills—a reporter—not write about it:

Do you think the game any advantage R

What game V

Checkers R

Amusement V

Don't you think people will make remarks R

Not if it is done by ourselves V

Yes have you any objections to Wills R

None if he does not publish it V

Wills was thoroughly impressed with the technology, calling it "another wonder of the age," according to Rogers. And so the telegraphers agreed that he could publish an account of the game, which perhaps was Vail's hope all along. The story was still being prepared for publication on 18 November when Vail tapped, "The Washington Chess Club challenge Baltimore to a game."

Vail's 1845 book about the telegraph includes a brief report on chess. He writes that in the Washington–Baltimore match, seven games were played, totaling 686 moves "transmitted without a single mistake or interruption." These details reappear in *The Book of the First American Chess Congress*, which called the Baltimore–Washington games the first telegraphic chess match.

How did the games actually unfold? While many of the rules and conventions

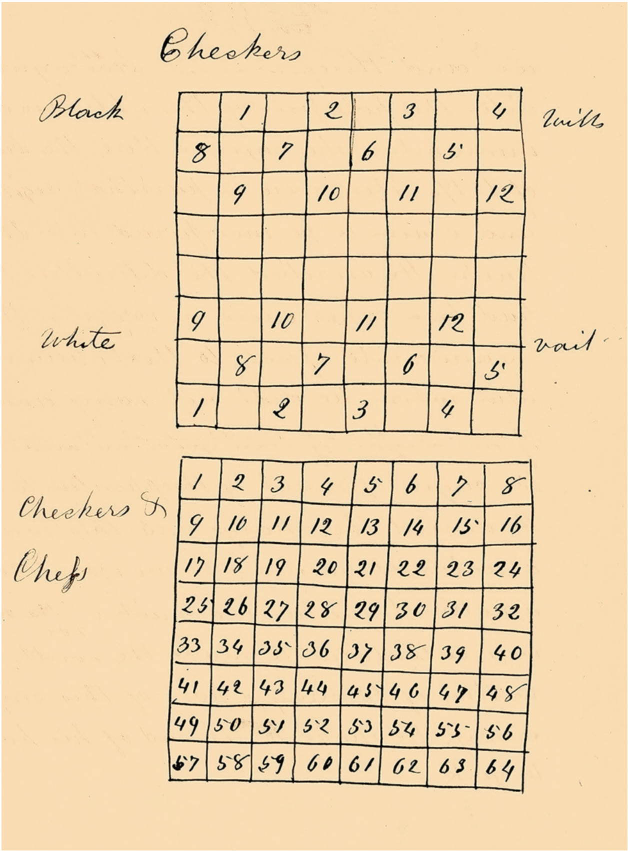

Checkers

Black | Vails

	1		2		3		4
8		7		6		5	
	9		10		11		12
9		10		11		12	
	8		7		6		5
1		2		3		4	

White | vail

Checkers & Cheſs

1	2	3	4	5	6	7	8
9	10	11	12	13	14	15	16
17	18	19	20	21	22	23	24
25	26	27	28	29	30	31	32
33	34	35	36	37	38	39	40
41	42	43	44	45	46	47	48
49	50	51	52	53	54	55	56
57	58	59	60	61	62	63	64

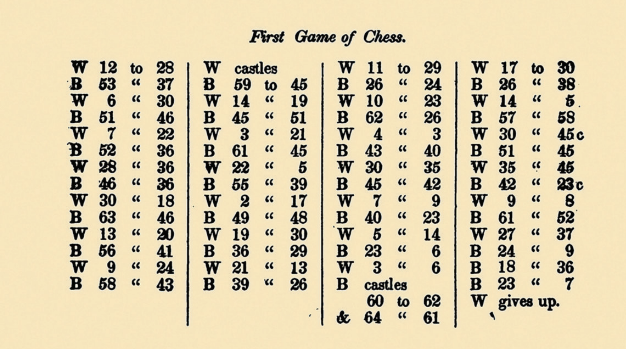

First Game of Chess.

W	**12**	to	**28**
B	**53**	“	**37**
W	**6**	“	**30**
B	**51**	“	**46**
W	**7**	“	**22**
B	**52**	“	**36**
W	**28**	“	**36**
B	**46**	“	**36**
W	**30**	“	**18**
B	**63**	“	**46**
W	**13**	“	**20**
B	**56**	“	**41**
W	**9**	“	**24**
B	**58**	“	**43**
W	castles		
B	**59**	to	**45**
W	**14**	“	**19**
B	**45**	“	**51**
W	**3**	“	**21**
B	**61**	“	**45**
W	**22**	“	**5**
B	**55**	“	**39**
W	**2**	“	**17**
B	**49**	“	**48**
W	**19**	“	**30**
B	**36**	“	**29**
W	**21**	“	**13**
B	**39**	“	**26**
W	**11**	to	**29**
B	**26**	“	**24**
W	**10**	“	**23**
B	**62**	“	**26**
W	**4**	“	**3**
B	**43**	“	**40**
W	**30**	“	**35**
B	**45**	“	**42**
W	**7**	“	**9**
B	**40**	“	**23**
W	**5**	“	**14**
B	**23**	“	**6**
W	**3**	“	**6**
B	castles		
	60	to	**62**
&	**64**	“	**61**
W	**17**	to	**30**
B	**26**	“	**38**
W	**14**	“	**5**
B	**57**	“	**58**
W	**30**	“	**45c**
B	**51**	“	**45**
W	**35**	“	**45**
B	**42**	“	**23c**
W	**9**	“	**8**
B	**61**	“	**52**
W	**27**	“	**37**
B	**24**	“	**9**
B	**18**	“	**36**
B	**23**	“	**7**
W	gives up.		

← To represent the positions on their telegraph chess board [second handwritten grid], Alfred Vail and Henry Rogers assigned a unique number to each of the 64 squares. In the first game of telegraph chess [bottom], White gave up.

of chess would be familiar to modern chess players, the games were unusual in other ways. At the time, the standard way of describing chess moves was descriptive notation, says John McCrary, a former president of the United States Chess Federation who maintains a collection of some 100 chess books from before 1900. For instance, “pawn to queen’s bishop’s four” described moving a pawn in front of the bishop on the queen’s side of the board to the fourth row from the bottom.

Before the electrical telegraph, such descriptions would have been used in correspondence chess, played by mail. And *The Oxford Companion to Chess* (1984) describes a proposed 1823 match between France and England that intended to use semaphore telegraph, although the notation was either never planned or has been lost to time.

But Vail and Rogers used a system that assigned a unique number to each of the 64 squares. So “pawn to queen’s bishop’s four” would have been rendered as “11 to 27.” Though the game itself can be remarkably complex, that system allowed individual moves to be communicated simply. “The exchange of information in chess is relatively low,” says David Kazdan, an engineer at Case Western Reserve University who has recently overseen a renewed collaboration between the school’s radio and chess clubs. “You don’t need much of a communication channel to play chess.”

Vail’s book logs the moves for two of the chess games, and both accounts include an illegal move—probably errors that were introduced later. The accounts in Vail’s telegraphic journal, on the other hand, appear accurate, and even include a real-time correction of one move.

In Vail’s journal, Washington claims the white pieces, but close examination shows that Washington either played the first move as black, or the board was mirrored left to right. At the time, the white pieces did not always take the first move. The sides agreed to a limit of 10 minutes per move, even though time controls weren’t common in chess, and the chess clock had not yet been invented. And while Vail wanted “first rate players,” McCrary calls the overall play weak, with a poor understanding of the long-term planning needed to coordinate all of the pieces. Both teams also made tactical errors. For example, in the second game, Washington overlooked that one of their pawns was overworked defending two other pieces simultaneously, watched as Baltimore captured one of the pieces, and elected not to retaliate in order to continue defending a more valuable knight. “Even with changing conventions of that time, what was there in the description was atypical,” says McCrary.

The teams took a break during the first game and then reconvened on 28 November. With a pawn in position to advance to the last row, where it would be “promoted”—that is, replaced by a more dangerous piece of the player’s choice—Baltimore swept in with its queen and readied checkmate in one move. Unable to salvage the game, Washington resigned. “Ha ha,” wrote Rogers. “Ha ha,” responded Vail.

There is no record of overall standings, and no winner was declared after all games had been played.

TOP: SMITHSONIAN INSTITUTION ARCHIVES; BOTTOM: GOOGLE BOOKS

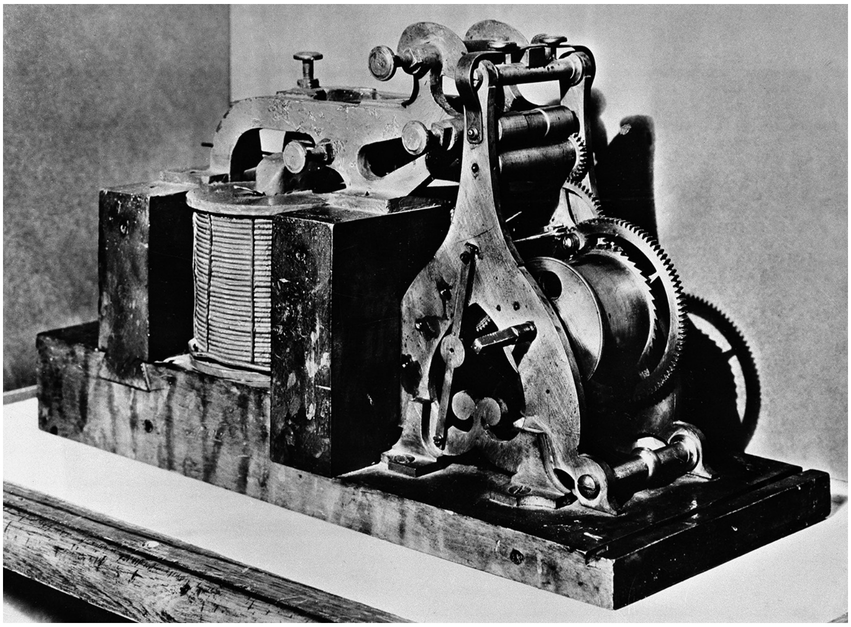

On 24 May 1844, this telegraph register received the first message sent by telegraph: “What hath God wrought.”

By today’s standards, the hardware that relayed the moves was simple, mainly consisting of a battery, a switch, and a magnet. “It’s not all that different from a doorbell,” says David Hochfelder, a historian at the State University of New York at Albany who has studied the early American telegraph.

Laying the line between the two cities had been difficult, with costly delays after failed attempts to bury the cable and to use cheaper noninsulated wire. Eventually, overhead insulated copper wire was strung the distance between poles.

Years before the chess match, Morse had considered a messaging system that used only numbers, which corresponded to set words or phrases listed in a code book. But he soon realized that a practical communications service would need an alphabetic component to spell out proper names.

This led to Morse’s eponymous code, which assigns a series of short and long signals to different alphanumeric characters. By tapping on a key, telegraph operators would interrupt a battery-powered current that ran the length of the telegraph wire. At the other end, an electromagnet moved a stylus, pen, or pencil, to mark a piece of paper with the corresponding dots and dashes, which an operator would then read. (The sounder, which turned the signals into audible sounds, hadn’t yet been invented.)

During the games, the telegraph operators occasionally asked each other how many people were in the room. At times, a dozen kibitzers looked on. At others, only the rotating cast of chess players and telegraph operators was present.

The Baltimore-Washington telegraph line was an immediate hit with a general public that embraced science through lectures and popular books

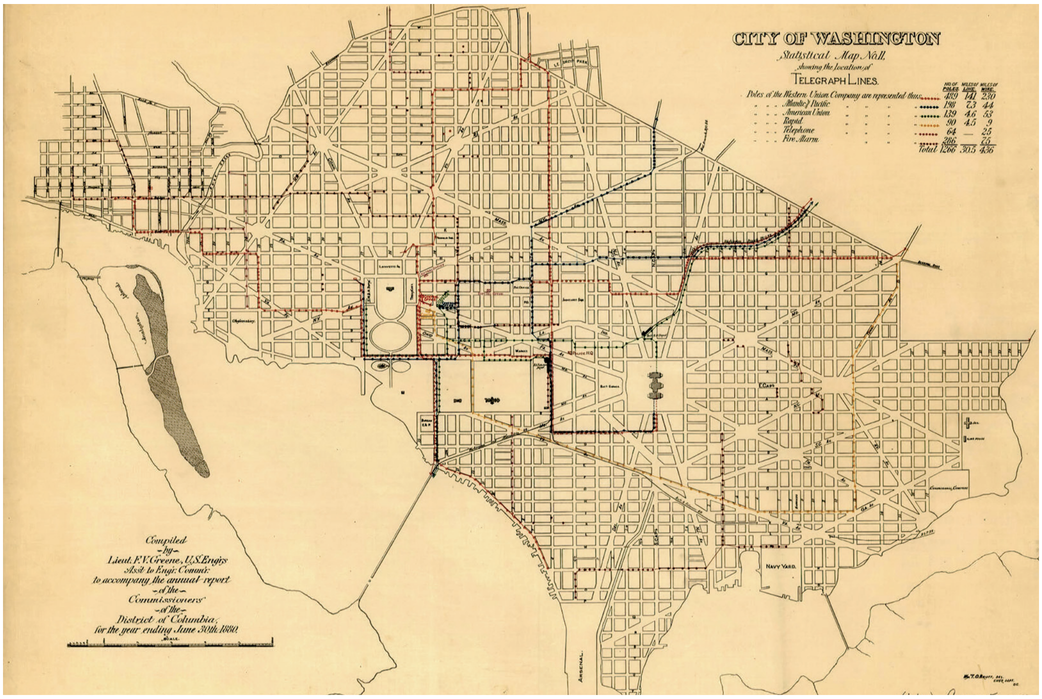
CITY OF WASHINGTON
Statistical Map No. II,
showing the location of
TELEGRAPH LINES.
NO. OF POLES. MILES OF LINE. MILES OF WIRE.
Poles of the Western Union Company are represented thus 489 14.1 23.0
Atlantic & Pacific 198 7.3 4.4
American Union 139 4.6 5.3
Rapid 90 4.5 .9
Telephone 64 — 2.5
Fire Alarm 286 7.5
Total 1266 30.5 43.6
Compiled
-by-
Lieut. F. V. Greene, U.S. Engrs.
Asst. to Engr. Commr.
to accompany the annual report
-of the-
Commissioners
-of the-
District of Columbia,
for the year ending June 30th 1880.
SCALE.
NAVY YARD
ARSENAL

The dotted lines in this 1880 map of Washington, D.C., reveal the telegraph network’s expansion following the opening of the Baltimore-Washington line in 1844.

TOP: AP; BOTTOM: LIBRARY OF CONGRESS

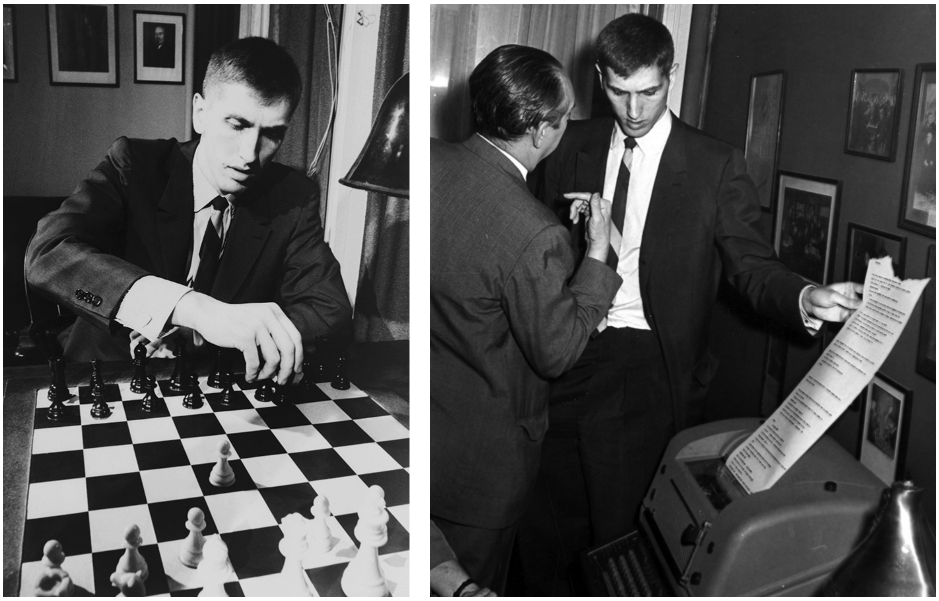

When U.S. chess grandmaster Bobby Fischer [left] was prevented from attending an international tournament in Havana, his moves were relayed via teletype.

LEFT: EVERETT COLLECTION HISTORICAL/ALAMY; RIGHT: SMITH ARCHIVE/ALAMY

and magazines. *Scientific American* was founded in 1845, for example. But people were more curious to see the telegraph at work than they were to use its services, even though the line operated free of charge for the first year. “Operators tended to show its capabilities rather than handling actual message traffic,” says Hochfelder.

The lack of activity is sometimes evident in the telegraph journal. Many of the messages are purely functional (“I am ready,” “stop 30 minutes”); simple greetings; notifications of letters sent and received; or requests for newspapers. The Baltimore end of the telegraph was in the Mt. Clare station of the B&O railroad, and the telegraph line ran alongside the tracks. Mail delivered by train took half a day door to door, says Hochfelder, and the telegraph offered little practical advantage.

On 5 December 1844, Rogers wrote to Vail:

“I hear from several sources that we are making rather an unfavorable impression with the religious part of the community, and I am under the impression if we continue after the present party is through that we will be injured more than any benefit might or can be derived from it.”

The exact nature of the religious community’s complaint with telegraph chess is unclear.

Although Morse wrote to Vail on the day of the first chess game that he “was much pleased with your game of drafts,” he came to feel that chess was too frivolous for the telegraph, as noted by the chess writer Tim Harding in his *Correspondence Chess in Britain and Ireland, 1824–1987* (McFarland, 2011). Whatever the reasons, it appears that after 17 December 1844, no more chess was played on the line. And in the end, Congress didn’t fund a telegraphic connection to New York, nor did it acquire perpetual rights to the telegraph, in part because Morse’s business partner had other designs, says Hochfelder. The Baltimore–Washington line operated under the auspices of the Postal Service from 1845 to 1847, when funding ended.

After that, the U.S. telegraph thrived in private ventures. Over the next few years, companies built lines and networks to connect cities across the country. Most notably, Ezra Cornell’s Western Union completed a transcontinental telegraph in 1861, and eventually became a monopoly. Ordinary people rarely used the telegraph, says Hochfelder, but it transformed finance and journalism.

Meanwhile, telegraph chess was taken up elsewhere. In 1845, for example, a match between London and Gosport, England, involved inventor Charles Wheatstone and chess master Howard Staunton. But it would take another few decades for telechess to become more widespread, with prominent club matches played over telegraph from the 1890s into the 1920s.

High-level chess competitions tend to be held in person, but such games have been played remotely from time to time. For example, in 1965, U.S. grandmaster Bobby Fischer relayed his moves by teletype over telephone lines from New York City to Havana, after the U.S. State Department prevented his attending a tournament there. And in 1999, a couple of years after losing a rematch to the IBM supercomputer Deep Blue, world champion Garry Kasparov played a promotional game against a team representing “the world,” which consulted on moves via a Microsoft forum.

Today, the internet has taken telecom chess to fabulous new heights, with one site alone, chess.com, often hosting up to 20 million games a day. Indeed, the growth in online play has sometimes stretched the capacities of the servers and the engineers who maintain them.

Why have technologists taken the opportunity to advance chess using so many generations of telecommunications? It may simply be that chess is popular, and by its nature can actually be played with short messages and perfect information, unlike soccer or poker.

But is there something more, maybe a natural affinity? “There are similarities in thinking processes [between] engineering design and the sort of puzzle solving that a chess game involves,” says Kazdan of Case Western Reserve. The connection may be one-sided, though. “Many engineers like chess. I’m not sure many chess players like engineering.” ■

CASE WESTERN RESERVE UNIVERSITY
Case School of Engineering

FACULTY POSITION IN DEPARTMENT OF ELECTRICAL, COMPUTER, AND SYSTEMS ENGINEERING

Case Western Reserve University, Cleveland, Ohio

The Department of Electrical, Computer, and Systems Engineering (ECSE) at Case Western Reserve University (CWRU) invites applications for one tenured faculty position at the Full Professor level. Appointments will be considered for starting dates as early as July 1, 2026. The successful candidate will be the inaugural holder of The Dr. Farrokh Khatibi Professorship endowed chair in the ECSE Department in the area of interdisciplinary advanced telecommunications. The candidate must have a Ph.D. degree in Electrical Engineering, or a closely related field, and be an internationally renowned thought leader with an outstanding record in scholarly research, fundraising, teaching excellence, and leadership in their professional community. The candidate will hold primary appointments in Case School of Engineering, although in many instances they will be eligible for additional appointments within other academic units of CWRU.

The Case School of Engineering is a recognized leader in innovative technologies such as robotics, artificial intelligence/machine learning (AI/ML), biomedical imaging, human-machine interaction, and precision diagnostics & therapeutics. The next wave of innovations in these fields will depend on the ability of devices, robots, and computational nodes to communicate wirelessly — rapidly and with high fidelity and security. The successful candidate is thus envisioned to lead transformative, interdisciplinary research initiatives that broadly focus on next-generation telecommunication technologies. The specific areas of interest to the department include, but are not limited to, **1)** design of high-performance analog/radio-frequency (RF) integrated circuitry for next-generation wireless connectivity, **2)** emerging applications of AI/ML in the design automation and optimization of analog/RF circuitry, **3)** signal processing for next-generation wireless telecommunication systems, and **4)** advanced wireless telecommunication systems and networks for AI/ML-enabled connectivity.

Applicants should provide the names and contact information of six references in addition to a cover letter, CV, research, teaching, and service/leadership statements. The search will continue until the position is filled; however, for maximum consideration the applicants should submit their materials by March 15, 2026. We welcome all nominations and applications for this available position; please submit your application materials to **https://apply.interfolio.com/178627**.

In compliance with the City of Cleveland's Pay Transparency Ordinance (effective October 27, 2025), the annual starting base salary range for this position is from $185,000 – $220,000. CWRU considers factors such as (but not limited to) scope and responsibilities of the position, candidate's work experience, education/training, key skills, internal peer equity, as well as market and organizational considerations when extending an offer. University employees enjoy a comprehensive benefits package that includes excellent healthcare, retirement plans, tuition assistance, and paid time off.

Case Western Reserve University strives to maintain a diverse and inclusive work environment. All applicants are protected under Federal law from discrimination based on race, color, religion, sex, national origin, disability, age, and genetics.

Case Western Reserve provides reasonable accommodations to applicants with disabilities. Applicants requiring a reasonable accommodation for any part of the application and hiring process should contact the Office of Equity by phone at 216.368.3066 or by email at **equity@case.edu**. Determinations as to granting reasonable accommodations for any applicant will be made on a case-by-case basis.

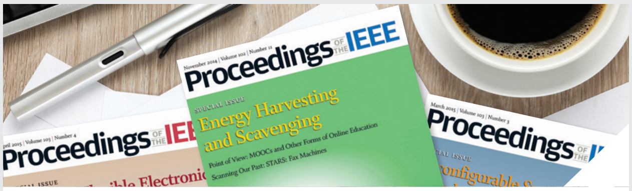

One of the most influential reference resources for engineers around the world.

For over 100 years, *Proceedings of the IEEE* has been the leading journal for engineers looking for in-depth tutorial, survey, and review coverage of the technical developments that shape our world.

To learn more and start your subscription today, visit **ieee.org/proceedings-subscribe**

IEEE

The Henry M. Rowan College of Engineering at Rowan University seeks a strategic, collaborative, and innovative leader to serve as the new **Head of the Department of Electrical and Computer Engineering.**

The new Head will spearhead the ECE Department to grow its undergraduate and graduate programs, drive innovation in education, research and scholarship, and further build upon its strong community of students and faculty.

For more information and to apply, visit go.rowan.edu/ECEDeptHeadSearch

RowanUniversity

IEEE
Foundation

Make it
Monthly

Join our community of monthly donors who compound their philanthropy for year-round impact.

Scan to Get Started

ieeefoundation.org/monthly

IEEE

BY PAUL JONES

The Attraction of Blackberries

The first time she tried to seduce me,
(atoms falling in a vacuum)
she asked about blackberries—
(every mass exerts some gravity)

did I know their season, where they grow?
(galvanometers, gravimeters)
I could answer both easily—
(tools to measure small attractions)

down the dirt road in September.
(devices that report, don't interfere)
She eagerly went there with me.
(variations in readings occur)

We ate more berries than we kept.
(electron exchange may explain this)
The sweet dark juice painted our lips.
(equilibrium then entropy)

ISTOCK

Paul Jones was inducted into the NC State Computer Science Alumni Hall of Fame in 2021 and has published two poetry collections with Redhawk Publications.

IEEE *Access*®

Multidisciplinary ⋮ Rapid Review ⋮ Open Access Journal

Harness the publishing power of IEEE *Access*®.

IEEE *Access* is a multidisciplinary open access journal offering high-quality peer review, with an expedited, binary review process of 4 to 6 weeks. As a journal published by IEEE, IEEE *Access* offers a trusted solution for authors like you to gain maximum exposure for your important research.

Explore the many benefits of IEEE *Access*:

- **Receive high-quality, rigorous peer review** in only 4 to 6 weeks
- **Reach millions of global users** through the IEEE *Xplore*® digital library by publishing open access
- **Submit multidisciplinary articles** that may not fit in narrowly focused journals
- **Obtain detailed feedback** on your research from highly experienced editors
- **Establish yourself as an industry pioneer** by contributing to trending, interdisciplinary topics in one of the many topical sections IEEE *Access* hosts
- **Present your research to the world quickly** since technological advancement is ever-changing
- **Take advantage of features** such as multimedia integration, usage and citation tracking, and more
- **Publish without a page limit** for $1,750 per article

⋮ Learn more at ieeeaccess.ieee.org

IEEE

BY ALLISON MARSH

Past Forward

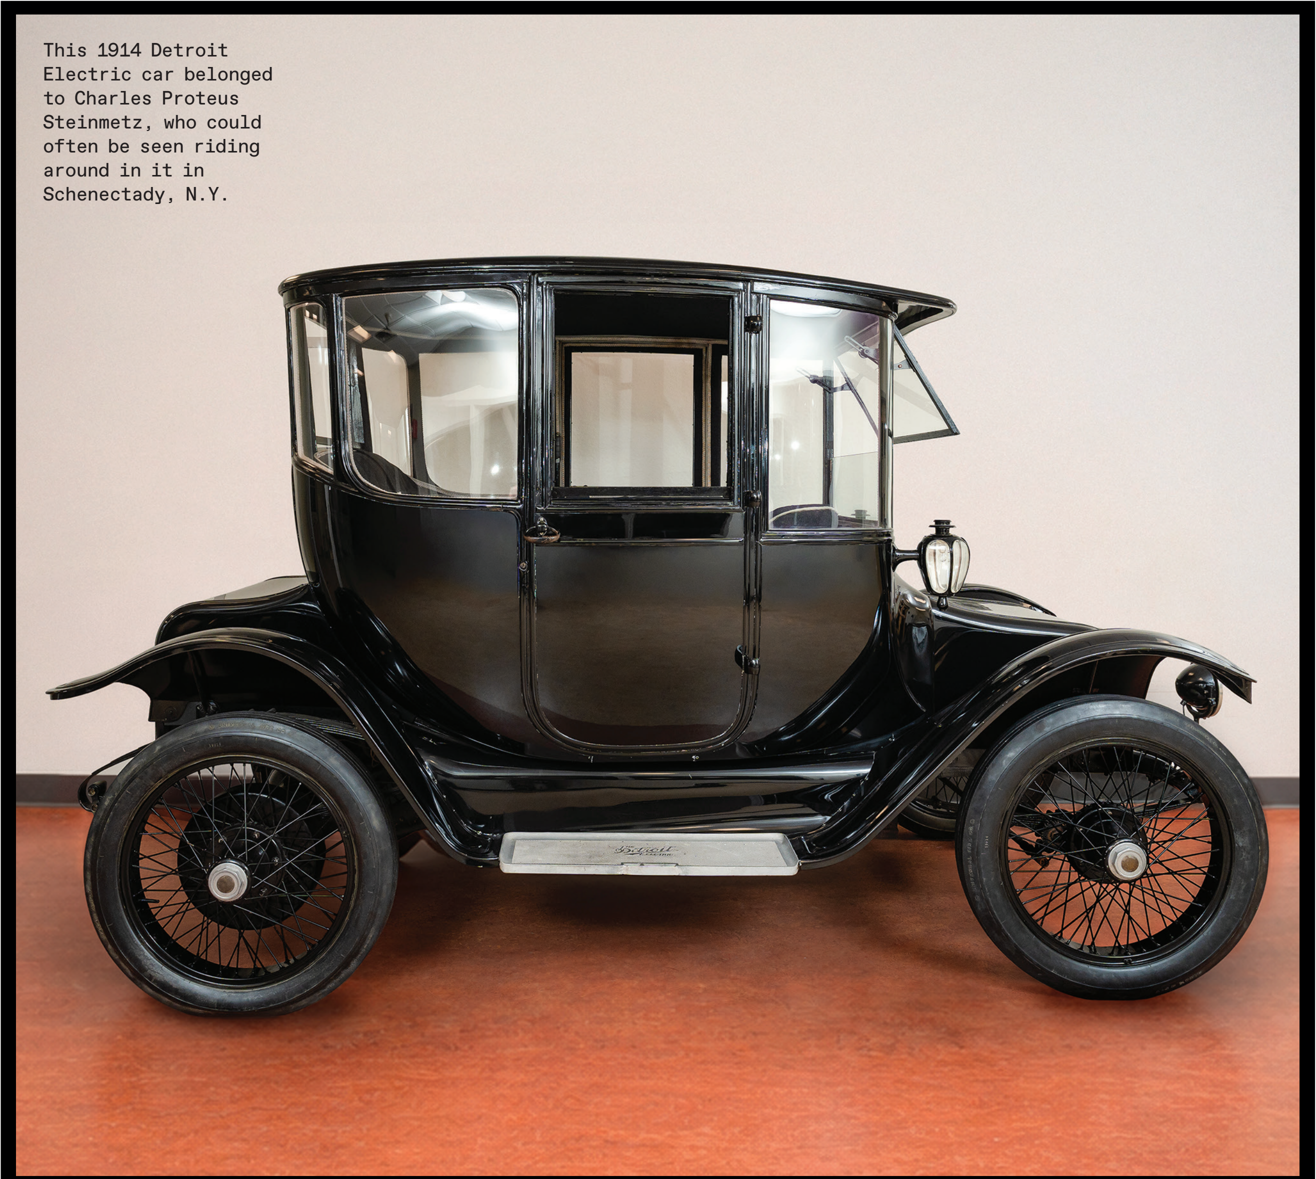

This 1914 Detroit Electric car belonged to Charles Proteus Steinmetz, who could often be seen riding around in it in Schenectady, N.Y.

Charles Proteus Steinmetz Loved His Electric Car

He was an intellectual titan of the electrification of the United States. Charles Proteus Steinmetz was also an evangelist for electric vehicles. On fine days, he could be seen decked out in a suit and top hat, smoking his trademark BlackStone panatela cigar and riding around Schenectady, N.Y., in his 1914 Detroit Electric sedan. Sometimes, both he and his chauffeur sat in the backseat—you could control the car from both the front and the rear—so that it would appear to be a driverless car.

In 1971, this car was purchased by Union College in Schenectady, where Steinmetz had founded, chaired, and taught in the department of electrical engineering. It had been discovered rotting in a field, so it needed some work. Over the next decade, faculty and engineering students restored the EV to its former glory. Still in running condition, it's now on permanent display at the college, a legacy to Steinmetz's genius and passion. ■

FOR MORE ON CHARLES PROTEUS STEINMETZ AND HIS ELECTRIC CAR, see spectrum.ieee.org/pastforward-mar2026

PAUL BUCKOWSKI/UNION COLLEGE

IEEE

YOU

JOINED THE GROUP

Member Discounts

YOU

HAVE EXCLUSIVE ADVANTAGES

Visit ieee.org/discounts to access your exclusive member savings.

MATLAB
FOR AI

Boost system design and simulation with explainable and scalable AI. With MATLAB and Simulink, you can easily train and deploy AI models.

mathworks.com/ai

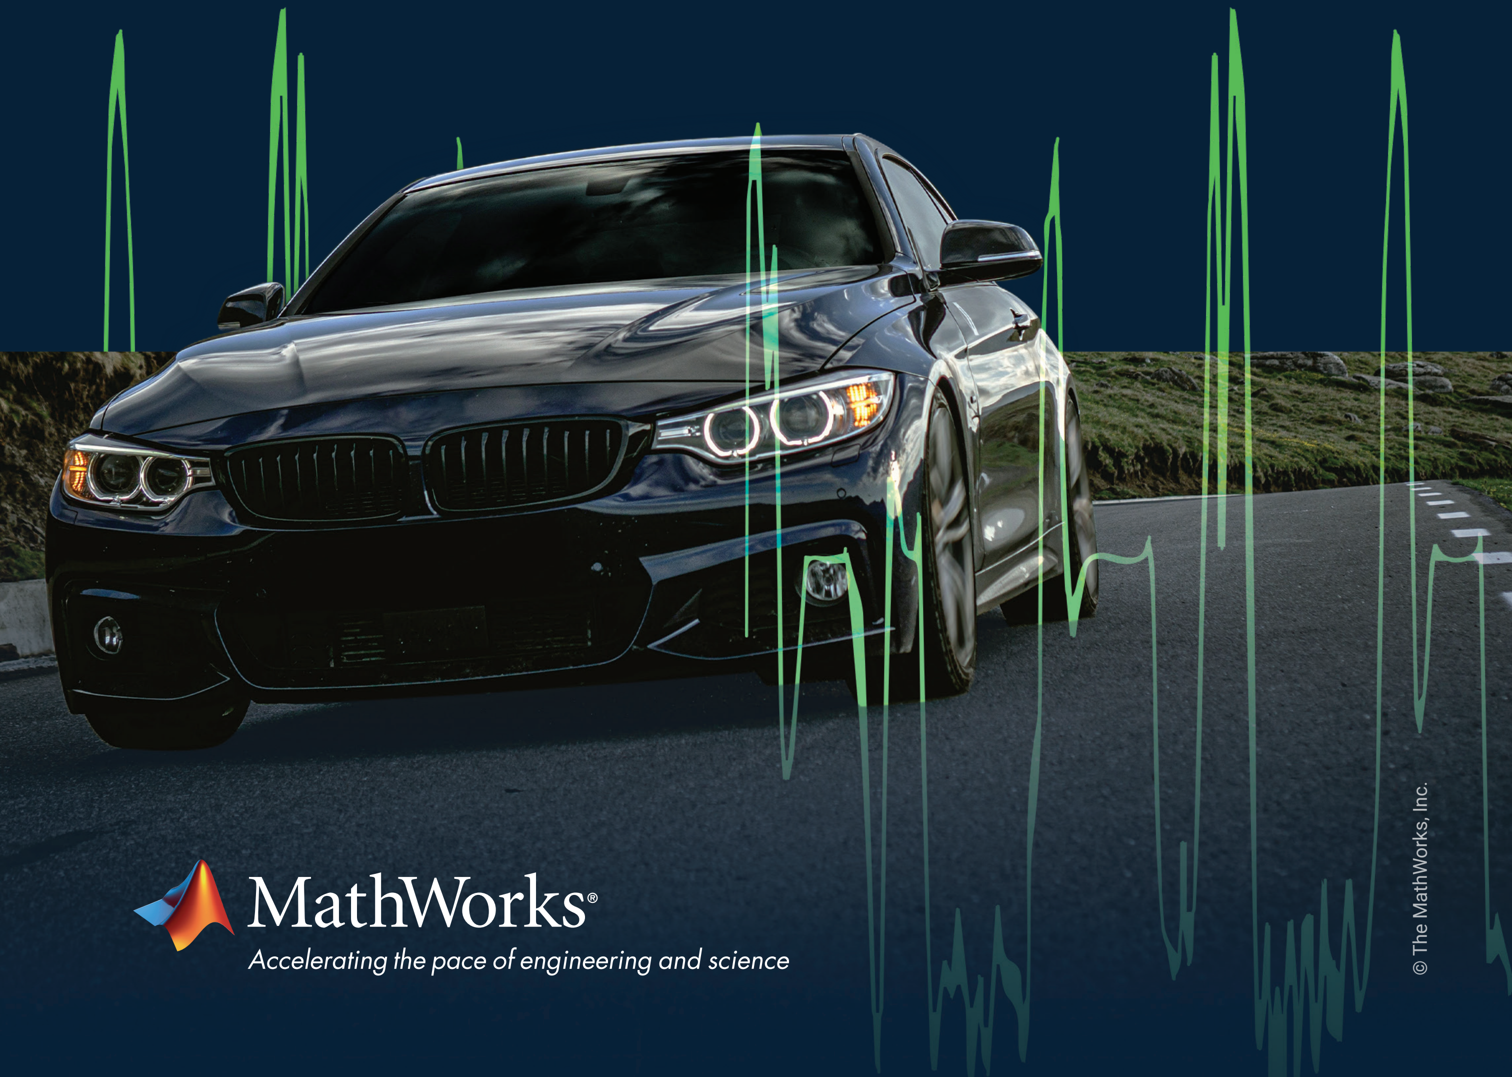

MathWorks®
Accelerating the pace of engineering and science

© The MathWorks, Inc.